# Real Consumption Measures for the Poorer Regions of the World

(preliminary draft for presentation
at the CRIW summer workshop - June 22, 2009)

Alwyn Young
Department of Economics
London School of Economics

# I. Introduction

Much of our current understanding of the factors behind growth and development, and our continuing attempts to deepen that understanding, is based upon crossnational estimates of levels and growth rates of real standards of living. Unfortunately, for many of the poorest regions of the world the underlying data supporting existing estimates of living standards is minimal or, in fact, nonexistent. Thus, for example, while the popular Penn World Tables purchasing power parity data set version 6.1 provides real income estimates for 45 sub-Saharan African countries, in 24 of those countries there has actually never been any benchmark study of prices.[1] In a similar vein, although the on-line United Nations National Accounts database provides GDP data in current and constant prices for 47 sub-Saharan countries for each year from 1991 to 2004, the UN statistical office which publishes these figures had, as of mid-2006, actually only received data for just under half of these 1410 observations and had, in fact, received no constant price data, whatsoever, on any year for 15 of the countries for which the complete 1991-2004 on-line time series are published.[2] Where official national data are available, fundamental problems of measurement produce a considerable amount of unquantified uncertainty. Thus, to cite one example, following the first national survey of the informal sector in Mozambique in 2004, the GDP estimates of nominal private consumption expenditure, extending all the way back into the

[1] See "Data Appendix for a Space-Time System of National Accounts: Penn World Table 6.1", February 2008. As explained in the source, expatriate post-allowance indices are used to extrapolate the price studies of benchmark countries to non-benchmark economies.

[2] This statement is based upon a purchase in 2006 of all the national accounts data records ever provided to the UN Statistics Division by member countries. When queried about the discrepancy between the completeness of their website and the data I had purchased, UN officials were quite frank about the difficulties imposed by the demands of users for a complete series, and their website openly explains that much of their data is drawn from other international organizations and extrapolations (http://unstats.un.org/unsd/snaama/metasearch.asp). Similar frankness concerning the need to use extrapolations from the data of other countries to fill in gaps is present on the World Bank data website (see http://go.worldbank.org/FZ43ELUKR0).

early 1990s, were doubled, producing a near doubling of GDP and, for any given benchmark price index, an equivalent doubling of measured living standards.

The paucity and poor quality of living standard data for less developed countries is well known and is motivating expanding efforts to improve the quality of information, as represented by the World Bank's International Comparison Programme and Living Standards Measurement studies. However, there already exists, at the present time, a large body of unexamined current and historical data on living standards in developing countries, collected over two decades as part of the Demographic and Health Survey (DHS). While the primary objective of these surveys has been to collect information on fertility and health, they have also, incidentally, developed a database of real material living conditions in the poorest regions of the world over the past two decades.

In this paper I use the DHS data to construct estimates of the level, growth rate and within region dispersion of real consumption in the rural and urban areas of 29 sub-Saharan and 27 other developing countries. These estimates have the virtue of being based upon a methodologically consistent source of information for a large sample of poor economies. As they make direct use of micro data, they incorporate an evaluation of the uncertainty introduced by sample sizes and the reduced information implied by correlation within clusters. Instead of attempting to measure total nominal consumption and marrying it to independently collected price indices, they employ a sampling of the real consumption of durable goods, housing, nutrition & health, and household time. Methodologically, they explicitly recognize the error involved in product sampling by, first, placing greater weight on products which have a stronger statistical correlation with real incomes and, second, discounting the observations in product groups which are strongly autocorrelated. In sum, I estimate the level, growth and inequality of living standards in the rural and urban regions of 56 developing economies using a methodologically consistent

information set and econometric techniques that take account of correlation within clusters and product groups in producing point estimates and their associated standard errors.

The principal result of this paper, when contrasted with conventional sources, is that the growth of average consumption, in both poor sub-Saharan and non sub-Saharan countries, is proceeding at about 3 percent per annum, i.e. about 2.5 times faster than the 1.1 to 1.3 percent indicated in cross-national datasets. Part of this discrepancy stems from units of measure: my estimate of the growth of living standards is the mean of ln real consumption, while conventional data sets report the ln of mean consumption. The difference between the two measures depends on the degree of within country inequality and, as I show, this is declining rapidly in the poorer countries of the world, particularly outside of sub-Saharan Africa. My methodology allows me to construct ln of the mean equivalent estimates of living standards and these show non sub-Saharan growth (at 2.3 percent) to be roughly on par with that reported in international data sets (1.7 to 1.9 percent). However, Sub-Saharan Africa's growth, measured either way, at about 3 percent per annum, is spectacularly above the .7 to .8 percent suggested by cross-national data sets. This discrepancy cannot be eliminated with econometric technique or definitions of income, and stems from the fundamental fact that the DHS data show a very rapid rise in real consumption measures for that region.

The development of aggregate consumption indices from consistent international micro-data allows a richer observation and breakdown of the factors influencing levels and trends than is possible with conventional national accounts aggregates. In this paper, in particular, I focus on the important role played by the urban-rural gap, which I find is a powerful determinant of productivity levels and overall within country inequality. The urban-rural gap is very large and varies extraordinarily little across countries. Consequently, much of what appears as "productivity differences" in levels accounting exercises (e.g. Hall & Jones 1999, Klenow &

Rodriguez-Clare 1997), really stems from differences in degrees of urbanization. Globally, however, the urban-rural gap is declining, and this contributes powerfully to the global decline in within country inequality, although there is also, interestingly, a downward trend in residual, after accounting for education, inequality within urban and rural regions.

I present my methodology and results in stages, allowing the reader to more easily absorb the different components that make up the approach and also establishing, I hope, that the basic results concerning growth are extremely robust. I begin, in section II, by describing the DHS data and the durable goods, housing, nutrition & health, and household time measures of real consumption. Section III then presents an introduction to my methodology, showing how aggregate consumption data provides information on the ratio of the growth rate to the cross national standard deviation of real living standards. Intuitively, the trend and cross national dispersion of real consumption is related, through the income elasticity of demand, to the trend and dispersion of real incomes so that the ratio of these two measures for real consumptions provides information, once one adjusts for product specific idiosyncratic effects, on the equivalent ratio for real incomes. Section IV implements this idea, highlighting the gross inconsistency between the DHS and the most popular of real living standards measures, the Penn World Tables. Put simply, the DHS data imply a ratio of growth to cross sectional dispersion on the order of 2.5 times that present in the PWT.

Section V continues by showing how the use of the micro correlation between educational attainment and consumption levels present in household datasets allows one to infer the income elasticity of demand for each of my real consumption measures, thereby allowing the separate estimation of both the growth rate and standard deviation of real living standards. When implemented using the DHS data in section VI, these methods indicate that the discrepancy between the DHS and the PWT lies in the growth rate of real consumption (as noted above) and

not its cross national dispersion, which is roughly equivalent in the two sources. In an attempt to move to a reconciliation of the two estimates, section VII notes how, conceptually, within country income inequality can contribute to a systematic difference between estimates of the mean of the ln, such as the DHS, and national accounts based reports of the ln of the mean, such as the PWT. To this end, Section IX expands my micro-data based methodology to allow the simultaneous estimation of within country levels, trends and inequality of real consumption. When implemented, in section X, these methods allow me to produce ln of the mean equivalent measures for the DHS which moves my estimates systematically closer to PWT in levels and, in growth rates, reestablishes a virtual equivalence between the estimated growth of non sub-Saharan countries, although a vast discrepancy between the DHS and PWT remains for sub-Saharan Africa. Section XI finishes by analyzing the role of urban-rural differences in determining levels and trends of productivity and inequality, as described above, and Section XII concludes.

## II. Demographic and Health Survey Data on Living Standards

The Demographic Health Survey and its predecessor the World Fertility Survey, both supported by the U.S. Agency for International Development, have conducted irregular, but in-depth, household level surveys of fertility and health in developing countries since the late-1970s. Over time the questions and topics in the surveys have evolved and their coverage has changed, with household and adult male question modules added to a central female module, whose coverage, in turn, has expanded from ever married women to all adult women. I take 1990 as my starting point, as from that point on virtually all surveys include a fairly consistent household module with data on household educational characteristics and material living conditions that are central to my approach. In all, I have access to 135 surveys covering 56 developing countries, as listed in Appendix I. The sample consists of about 0.5 million households in sub-Saharan Africa and 1.1 million households in the rest of the world, including useful information on 3.1 million youths aged 6-24, 1.2 million currently married women aged 15-49, and 0.65 million children less than 3 years of age.[3]

The raw data files of the DHS surveys are distributed as standardized "recode" files. Unfortunately, this standardization and recoding has been performed, over the years, by different individuals using diverse methodologies and making their own, idiosyncratic, errors. This produces senseless variation across surveys as, to cite two examples, individuals with the same educational attainment are coded as having dramatically different years of education or individuals

[3] These numbers represent the sample with relevant information, not the the total number of individuals in the surveys. Thus, the 1.6 million households actually contain 3.5 million youths aged 6-24, 2 million women aged 15-49, and .9 million children aged 3 or less. Youth school attendance information is not collected in some surveys and, depending upon the survey, women are interviewed in depth (providing information on themselves and their young children) according to whether they have ever been married, or slept in the house the previous night or are usual members.

who were not asked education attendance questions are coded, in some surveys only, as not attending. In addition, there are underlying differences in the coverage of the surveys (e.g. children less than 5 years vs. children less than 3 years) and the phrasing and number of questions on particular topics (e.g. employment) which produce further variation. Working with the original questionnaires and supplementary uncoded raw data generously provided by DHS programmers, I have recoded all of the individual educational data, corrected a number of major coding errors in individual items, recoded variables to standardized definitions and, as necessary, restricted the coverage to a consistent sample (e.g. married women, children less than 3 years) and removed surveys with inconsistent question formats (in particular, regarding labour force participation). Appendix I lists some of the details.

I use the DHS data to derive 26 measures of real consumption distributed across four areas: (1) ownership of durables; (2) housing conditions; (3) children's nutrition and health; and (4) household time and family economics. Table I below details the individual variables and sample means. All of these variables are related to household demand and expenditure, broadly construed, and, as shown later, are significantly correlated with real household incomes, as measured by average adult educational attainment. I have selected these variables on the basis of their availability and with an eye to providing a sampling of consumption expenditures that would, through material durables, nutrition & health and household time, cover most of the budget of households in the developing world. I have made the decision to break measures of household time into different age groups to account for different demand patterns at different ages as the possibilities of substitution between home production, human capital accumulation and market labour evolve. Thus, for example, in richer households young women are more likely to be in school and less likely to be working in the late schooling years (ages 15-24), but, consequently, are more likely to be working as young adults (ages 25-49). Although males are included in

| Table I: DHS Real Living Standard Measures by Category | | | | | |
|---|---|---|---|---|---|
| | N | Mean | | N | Mean |
| Ownership of Durables | | | Housing Conditions | | |
| Radio | 1557550 | .574 | Electricity | 1534362 | .528 |
| Television | 1577616 | .405 | Tap Drinking Water | 1569114 | .451 |
| Refrigerator | 1473490 | .249 | Flush Toilet | 1449330 | .322 |
| Bicycle | 1489805 | .296 | Constructed Floor | 1400359 | .598 |
| Motorcycle | 1431210 | .102 | ln # Sleeping Rooms per Person | 717178 | -.927 |
| Car | 1460012 | .066 | | | |
| Telephone | 1130847 | .172 | | | |
| Children's Nutrition and Health | | | Household Time and Family Economics | | |
| ln Weight (100g) | 465085 | 4.44 | Attendifbng School (age 6-14) | 1916473 | .712 |
| ln Height (mm) | 454582 | 6.59 | Attending School (age 15-24) | 1219551 | .304 |
| No Diarrhea | 590540 | .799 | Working (women age 15-24) | 195060 | .416 |
| No Fever | 578304 | .676 | Working (women age 25-49) | 588049 | .554 |
| No Cough | 582544 | .658 | Gave Birth Past Year (age 15-24) | 289763 | .312 |
| Alive | 649386 | .930 | Gave Birth Past Year (age 25-49) | 898526 | .141 |
| | | | Ever Married (women age 15-24) | 726630 | .431 |
| | | | Ever Married (women age 25-49) | 1083877 | .936 |

Notes: All variables, other than ln weight, height and rooms per capita, coded as 0/1. Ownership of Durables: at least one such item in the household; Housing Conditions: constructed floor means made of other than dirt, sand or dung. Household Time: individual variables, i.e. coded separately for each individual of that age in the household; recent fertility and market participation refer to currently married women only. Children's Health: individually coded for each child born within 35 months of the survey; diarrhea, cough and fever refering to the absence of these for the individual in question (if alive) in the preceding two weeks; ln weight and ln height refering to measurements of living children at the time of the survey.

the schooling and children's health variables, I do not include separate time allocation measures for adult males because male questionnaire modules are less consistently available and male participation behavior, when recorded, is less strongly related to household income and, hence, by my methodology, would play little role in estimating relative income.

My approach will be to use the correlation between real consumption in a sample of products and relative household incomes, as measured by adult educational attainment, to draw inferences about levels and trends in relative regional incomes. While the full methodology is discussed in later sections, a few obvious concerns should be noted at this time. First, there is

likely to be a significant covariance, independent of income, between many of these variables at the national or regional level, so that they cannot be construed as a true, independently drawn, random sample of consumption products. The use of random effects within the four broad product groups (e.g. housing), in both trends and within regions, will address this issue, discounting the number of observations to the degree that there are strong empirical correlations within but not across groups, i.e. to the degree that the product groups represent correlated demand along narrow dimensions. Second, there are likely to be regional, idiosyncratic, factors affecting the measured levels of many of these variables, independent of relative national income. Again, these can be explicitly recognized with random effects at the product and product group level, identifying the degree to which there is correlation in levels within products and product groups that do not extend across all products, producing more efficient estimates of overall country real consumption levels with appropriate standard errors. Third, local, district level, infrastructure is an important determinant of the realized household consumption of some products. Arguably, individuals choose their residence precisely to get access to such communal infrastructure, and pay for it implicitly through the local cost of housing and land, so it should be viewed as an element of demand. Nevertheless, I will address this issue by estimating demand equations using cluster level random and fixed effects. The results are not dramatically different. Finally, a skeptic might question whether the consumption of some of these "products" is even related to real income. As presented later, my methodology is inherently "idiot proof", as products whose micro level correlation with real incomes is poor will play an insignificant role in determining the estimated growth and levels of regional incomes.

As a final comment on the data, I should note that the DHS codes households as living in urban (cities and towns) or rural (countryside) areas. In what follows, I estimate average levels

of demand and income at the urban and rural level. To calculate and report national averages, I use the average share, across all surveys for a particular country, of the urban and rural households in the total household weights given by the DHS.[4] For the 56 countries in my sample, the urban share varies from a minimum of .11 in Rwanda to .73 in Gabon, with a mean of .38 and standard deviation of .17. As shown later, differences in the levels and growth rates of urban and rural living standards, and their correlation with the degree of urbanization, play an important role in explaining the level and trends of national living standards and income inequality.

[4] I intend to expand the estimates to include a time series on the urban/rural share but, pending a more detailed examination of the DHS estimates of the degree of urbanization and how these compare to alternative sources, I discount any time series variation at this time and use the simple DHS average.

## III. Methods I: Using Consumption Aggregates

I introduce my methodology in stages: first, in this section, showing how the ratio of the growth to the standard deviation of living standards can be inferred from panel data on real consumption aggregates, then, in a later section, showing how the within survey micro-level correlation between consumption and educational attainment can be exploited to produce separate point estimates of both the growth and standard deviation of living standards and, yet further on, using the consumption correlation within households to infer the degree of within region inequality. Proceeding in this sequential fashion makes the pieces that make up the overall methodology transparent and digestible and, as the reader will see, highlights the fundamental inconsistency between the DHS data and the most recognized of international measures of living standards, the Penn World Tables.

Let some measure of the demand for product p be given by:

$$(1)\quad \ln(Q_p) = \alpha_p + \eta_p \ln(Y^N) + \vec{\xi}_p^{\,\prime} \ln(\vec{P})$$

where $\alpha_p$ is a constant, $\eta_p$ the quasi-income elasticity of demand, $Y^N$ nominal income, $\vec{\xi}_p$ a vector of own and cross quasi-price elasticities of demand, and $\ln(\vec{P})$ the associated vector of prices relative to some base. I use the term quasi in describing the elasticities, because Q need not be actual quantity demanded, but only some measure related to that quantity, such as the index in a probability model or an outcome of food demand such as body weight. Homogeneity of demand of degree 0 in income and prices implies that that the quasi-income elasticity of demand equals the negative of the sum of the own & cross quasi-price elasticities:

$$(2)\quad \eta_p = -\sum_q \xi_{pq}$$

Equation (2) holds even when Q is not strictly speaking quantity demanded, as anything associated with that demand should, equally, have the same homogeneity of degree 0 property.

To reformulate (1) in terms of real income, we add and subtract from nominal income the expenditure share weighted movement of prices from the base to produce

$$(3) \quad \ln(Q_p) = \alpha_p + \eta_p[\ln(Y^N) - \vec{\Theta}' \ln(\vec{P})] + \eta_p[\vec{\Theta}' - \vec{\xi}_p{}'/\eta_p]\ln(\vec{P})$$

where $\vec{\theta}$ is a vector of product expenditure shares.[5] I operationalize (3) empirically by taking the last term on the right hand side as the error term:

$$(4) \quad \ln(Q_p) = \alpha_p + \eta_p \ln(Y^R) + \eta_p \varepsilon$$

Clearly, $\vec{\Theta}$ and $\vec{\xi}_p/\eta_p$ are vectors whose components each sum to one and, when differenced, sum to zero. Consequently, uniform inflation drops out of the error term which, when normalized by the quasi-income elasticity, is a zero-weight average of relative price changes; something that, arguably, is homoskedastic across products.

Imagine real income is growing, worldwide, at a common rate, so that real income per capita in country c at time t is given by $\ln(Y^R_{ct}) = \ln(Y^R_c) + gt$, and let $\sigma_{Y^R}$ denote the standard deviation of ln permanent relative real income, $\ln(Y^R_c)$. Treating each country as a representative consumer, consider running for a single product p the random effects panel regression:

$$(5) \quad \ln(Q_{pct}) = \hat{a}_p + \hat{g}_p t + u_{pc} + e_{pct}$$

If ε in (4) is truly iid, then

$$(6) \quad \frac{\hat{g}_p}{\hat{\sigma}_u} = \frac{\eta_p g}{\eta_p \sigma_{Y^R}} = \frac{g}{\sigma_{Y^R}}$$

[5] These are actual product expenditure shares, and not quasi in any way, but, as will be seen, there is no need to actually ever compute them.

so that the time series/cross-sectional variation of average consumption levels gives information on the ratio of the growth rate to the underlying standard deviation of real incomes.[6] Unfortunately, $\varepsilon$ is not likely to be iid, as there are persistent influences on demand within a country other than real income, most notably sustained relative price differences. Hence, the underlying data are actually produced by the process:

$$(4)' \quad \ln(Q_{pct}) = \alpha_p + \eta_p \ln(Y^R_{ct}) + \eta_p \gamma_{pc} + \eta_p \varepsilon_{pct}$$

where $\gamma_{pc}$ represents a persistent country x product error and $\varepsilon_{pct}$ an iid residual error. In this case, (5) and (6), run on a single product, actually estimate

$$(6)' \quad \frac{\hat{g}_p}{\hat{\sigma}_u} = \frac{g}{\sigma_{Y^R+\gamma}}$$

It seems likely that there is a fair amount of independent variation in $\gamma_{pc}$ so that, on average, (6)′ will understate the ratio of growth to the standard deviation of real income.

The problem of persistent, non-income, country specific influences on the demand for indi-

[6] As all of this, and everything to come, is maximium likelihood, all statements about coefficients being equal to parameters are actually statements about asymptotic consistency. I generally have about 130 observations at the product level regression presented so far, and close to 3200 when I combine all products to produce my final estimates (discussed further below). For the micro-level analysis presented in later sections, there are typically at least 1/2 million observations, with thousands of observations per regional dummy.

I should note that (5) can also be run as a fixed effects model. In this case, however, one cannot simply calculate the standard deviation of the country fixed effects, as it is exaggerated by estimation error. Instead, one has to run the estimated fixed effects on a constant, explicitly taking into account the estimated variance of the first stage estimates, and using the estimated residual variance of this second stage to calculate the standard deviation of underlying real incomes

$$\text{i.e.} \quad \hat{u}_c = c + \varepsilon_u \quad \text{with } E(\vec{\varepsilon}_u \vec{\varepsilon}_u{}') = I(\sigma_u^2) + \Sigma(\hat{u}_c)$$

where I is a diagonal matrix and $\Sigma$ is the first step covariance matrix. As this procedure is more complicated, I relegate it to an appendix (available upon request from the author), where I show it produces virtually identical results. This is not surprising, as the assumption that the random effects in (5) are orthogonal to the constant and time trend is quite reasonable.

vidual products can be solved by extending the single product random effects regression to a random sample of products:

$$(7)\quad \ln(Q_{pct}) = \hat{a}_p + \hat{b}_p(\hat{g}t + u_c + e_{pc} + e_{pct})$$

where all error terms are uncorrelated with each other and across subscript categories, e.g. $E(e_{ic}e_{jc}) = 0$ ($i \neq j$). As it stands, (7) is not identified, so one normalizes $\hat{\sigma}_u = 1$, and estimates a vector of normalized constants $\hat{\mathbf{a}}$ and quasi-income elasticities $\hat{\mathbf{b}}$, normalized standard deviations of persistent product x country shocks $\hat{\sigma}_{pc}$ and residual iid shocks $\hat{\sigma}_{pct}$, and, of course, the growth rate $\hat{g}$. This last item, if the underlying data are produced by (4)′, actually estimates the ratio of real growth to the standard deviation of real income:

$$(8)\quad \hat{g} = \frac{g}{\sigma_Y}$$

It is obvious that there are various generalized least squares extensions of (7) that will improve the efficiency of the estimates and allow them to more accurately reflect the true informational content of the sample. Thus, for example, there is likely to be significant cross country correlation in product specific growth rates, as relative prices, globally, follow particular trends. This is incorporated by extending the random effects framework to:

$$(9)\quad \ln(Q_{pct}) = \hat{a}_p + \hat{b}_p(\hat{g}t + e_pt + u_c + e_{pc} + e_{pct})$$

where $e_p$ represents correlation in product specific trends. Similarly, the individual products may not be a true random sample, so that there are significant non-income related correlations in levels within countries and in growth rates across countries:

$$(10)\quad \ln(Q_{pct}) = \hat{a}_p + \hat{b}_p(\hat{g}t + e_pt + e_Gt + u_c + e_{pc} + e_{Gc} + e_{pct})$$

where G denotes a product group (e.g. the housing data in Table I earlier), $e_G$ captures commonalities in product group growth rates and $e_{Gc}$ commonalities in country specific product group consumption levels.[7]

Since the estimated regression coefficients are simple constants and time trends, the additional random effects in (7), (9) and (10) have little to do with the traditional weighting of "between" and "within" estimators. Instead, large estimated values of $\hat{\sigma}_{pc}$, $\hat{\sigma}_{Gc}$, $\hat{\sigma}_p$ and $\hat{\sigma}_G$ indicate that there is considerable correlation between the error terms, i.e. that the nominal number of observations considerably overstates the true information in the sample, producing larger standard errors for the estimated parameters. This also influences, more subtly, the coefficient estimate of $\hat{g}$. As the informational content of the sample is reduced, the estimate of the true variation in country living standards ($\hat{\sigma}_u$), for any given observed differences in consumption levels, falls.[8] As $\hat{\sigma}_u$ is normalized to equal 1, this raises the coefficient estimate of $\hat{g}$, the ratio of growth to cross sectional variation.

---

[7] As before, all shocks are orthogonal to each other. Thus, for example, $e_{pc}$ represents the residual country x product correlation after accounting for country x product group correlation $e_{Gc}$. All of these errors are normalized by the income elasticity, for the reason given at the beginning of this section. However, I've experimented with assuming that this is not the case for the residual iid shock $e_{pct}$ and have gotten very similar estimates of g (presented in an appendix, available upon request from the author).

[8] This is most easily seen by thinking of the two step fixed effects procedure described in an earlier footnote, where one first uses an equation such as (10) to estimate country level fixed effects and then projects these on a constant, taking into account the estimated covariance matrix of the fixed effects from the first step. A higher observed correlation within products raises the standard errors of the estimated fixed effects implying a lower estimate of the true, underlying variation in the second step. As a simple example, consider the case where each product group is made up of 10 duplicate observations of the same product. In this case, the introduction of first step product group random effects has no effect, whatsoever, on the point estimates of the country level fixed effects, but raises their standard errors. Any given observed mean variation in country consumption levels across all products is now associated with more error (there being fewer real observations), and hence more likely to be the outcome of sampling variation rather than true underlying differences.

As a final refinement, I should note that efficiency requires that one explicitly take into account the fact that the observations of the left-hand side variable, i.e. consumption levels, are estimated from data in a preliminary first-step. To this end, the covariance matrix for the likelihood should be augmented with the covariance matrix of the first step estimates, i.e. $\Sigma(\text{second step}) = \Sigma(\text{RE}) + \Sigma(\text{first step})$, where the first term is the covariance matrix described by the random effects model (e.g. eqn. (10)) and the second is the covariance matrix of the left-hand side variable, as estimated in the first step. A larger estimated first step covariance matrix compresses the inferred true, underlying, variation in the sample. For estimation based upon national or urban/rural averages calculated from individual surveys, in the following section, this is not of great import, as the sample sizes ensure that the standard errors are miniscule. Further on in the paper, however, where micro-level correlations are used to estimate the quasi-income elasticities, this procedure places less weight on products whose association with income is weaker, and hence provide less information about trends and relative levels of real income.[9]

To summarize, under the assumption of a ln linear income elasticity of sorts, the time series and cross-sectional variation in a measure of the real consumption of a product provides information on the ratio of the growth to the standard deviation of real income. Econometrically, estimation of this ratio can explicitly account, in a variety of ways, for correlation within products and product groups in both within country levels and cross-national growth rates, producing standard errors that properly reflect the amount of independent information in the product sample. The use of estimated means as a left hand side variable requires that the first stage

[9] I should note that the standard errors of the second-step estimated coefficients and random effects standard deviations should also be modified to take into account their dependence on the first-step parameter estimates and their standard errors, i.e. a two-step calculation of standard errors. However, given the complexity of some of the likelihoods, the calculations are quite difficult and time consuming. Hence, having checked a few cases and found that this does not substantially alter the results, I leave this particular refinement for a later draft.

covariance matrix be incorporated in the second stage likelihood functions description of the covariance matrix. These econometric refinements play a significant role later on, when I separately estimate the trend and standard deviation of real income. They are not, however, necessary to establish the central result of this paper, as presented in the next section, i.e. that there is a glaring inconsistency between the relative degree of time series and cross-sectional variation in real consumption present in the DHS and the PWT.

## IV. Results I: The Gross Inconsistency Between the DHS and PWT

Table II below presents product-level estimates of the ratio of the growth to the standard deviation of real income following equations (5) and (6)′ described in the previous section. For each entry in the table, the dependent variable is a measure of the mean consumption of a product in a particular survey, i.e. a country x time panel for a given product. For ln rooms per capita and children's ln weight and height, this is the sample mean. For the remaining dichotomous variables, coded in the surveys as 0 or 1, the measure used in the table is the logit of the mean value, i.e. $\ln(\bar{x}/(1-\bar{x}))$.[10] I calculate these values at the urban and rural level for each survey and aggregate them using the mean urban/rural weight to produce a national measure.[11] The estimated variance of these first step estimates is adjusted for clustering and then additively incorporated into the likelihood of the second step random effects model, as described earlier above.[12] Although these refinements matter in later sections, in this table they are not crucial, and estimates calculated using simple national means (without urban/rural weighting), with or without adjustment for the first step variance or clustering, and even including more complicated

[10] I use the logit as my baseline, rather than simply the sample mean value, because later in the paper I will need to estimate probability equations at the micro level, i.e. deal explicitly with the fact that the variable can only take two values, and the logit is a convenient probability model that is easily extended to random and fixed effects specifications. However, as shown in this section and later in the paper, the results do not depend upon the choice of functional form.

[11] As for the independent variable "time", since the surveys are executed over a period of months, I code each survey as taking place at the average date at which the households were interviewed, with each month coded as 1/12 of a year. Section II and Appendix I provide further details on the definition and construction of the variables.

[12] To clarify, the first step predicted values are equal to those one would get if one ran, for each urban/rural region in each survey, a linear regression or logit model on a constant. The standard errors of these urban/rural means can then be adjusted for clustering using the usual robust "sandwich" estimator of variance. The national mean for each survey is the weighted average of the urban/rural means, and its variance is the sum of the square of urban/rural weights times the individual urban/rural variances.

Table II: Product Level Estimates of Growth/Standard Deviation
Dependent variable = urban/rural weighted country means
$y_{pct} = a_p + g*t + u_{pc} + e_{pct}$, reporting $\hat{g}/\hat{\sigma}_u$

| Durables | | Housing | |
|---|---|---|---|
| Radio | .054 (.011) | Electricity | .031 (.004) |
| Television | .043 (.005) | Tap Water | .001 (.006) |
| Refrigerator | .025 (.004) | Flush Toilet | .031 (.006) |
| Bicycle | .039 (.006) | Constructed Floor | .016 (.004) |
| Motorcycle | .029 (.006) | ln(Rooms/Capita) | .017 (.008) |
| Car | .015 (.005) | | |
| Telephone | .065 (.012) | | |
| **Children's Health** | | **Family Economics** | |
| ln Weight | .028 (.007) | At School (6-14) | .055 (.009) |
| ln Height | .045 (.010) | At School (15-24) | .045 (.010) |
| No Diarrhea | .032 (.014) | Working (15-24) | .019 (.011) |
| No Fever | .061 (.017) | Working (25-49) | .033 (.013) |
| No Cough | .047 (.017) | Birth (15-24) | .051 (.013) |
| Alive | .054 (.008) | Birth (25-49) | .033 (.006) |
| | | Marriage (15-24) | .015 (.005) |
| | | Marriage (25-49) | .014 (.004) |

random effects (at the cluster level) estimation of regional means, are all virtually identical. To avoid pointless repetition, these are placed in an appendix, available upon request from the author. Across all 26 products, the average of the growth to standard deviation ratio and its standard error is .034 (.008).

Table III runs the same random effects regression (5) used in Table II on the real consumption data of the Penn World Tables. In this case, both the growth and standard deviation are identified, as there is no income elasticity that needs to be implicitly estimated, but only their ratio can meaningfully be compared to the DHS results presented up to this point in the paper. I use as my observations the 132 PWT country x year observations corresponding to the countries

| Table III: PWT & UN Based Estimates of the Growth and Standard Deviation of Real Living Standards $y_{ct} = a + g*t + u_c + e_{ct}$ | | | | | | |
|---|---|---|---|---|---|---|
| | Penn World Tables 6.2 Per Capita | | Penn World Tables 6.2 Per Equivalent Adult | | UN National Accounts Per Capita | |
| | GDP | Private Consumption | GDP | Private Consumption | GDP | Private Consumption |
| $g$ | .016 (.002) | .013 (.002) | .014 (.002) | .011 (.002) | .016 (.002) | .013 (.002) |
| $\sigma_u$ | .730 (.070) | .667 (.064) | .705 (.067) | .642 (.061) | .764 (.073) | .699 (.067) |
| $\sigma_{ct}$ | .092 (.007) | .097 (.008) | .091 (.007) | .097 (.008) | .117 (.009) | .105 (.008) |
| $g/\sigma_u$ | .021 (.004) | .020 (.004) | .020 (.004) | .018 (.004) | .022 (.004) | .019 (.004) |

Notes: Calculated using PWT Laspeyres measures of GDP. PWT chain measures produce identical results. PWT calculates equivalent adults by assigning a weight of .5 to persons under 15. UN measures are in constant market exchange US dollars, with ad hoc PPP adjustments (see text).

and dates of my DHS surveys.[13] As the dependent variable I use per capita and per equivalent adult measures of real gdp and real consumption. As the reader can see, the PWT ratio of growth to standard deviation for these measures varies from about .021 down to .018. This is around 1/2 to 2/3 the average present at the product level in the DHS data, as shown in Table II, where 3/4 of the products show ratios higher than .021. It is immediately apparent that there is a vast discrepancy between the degree of growth to cross sectional variation present in the PWT and that present in the DHS. As a cross-check, Table III runs the same analysis on the UN National

[13] I average/weight the PWT data for the years in which each DHS survey takes place (e.g. 2003-2004) based upon the date in which the average household was surveyed. For the countries in my sample, the PWT 6.2 data end, mostly, in 2003 and 2004. For 38 of my DHS surveys the survey date is at or past the last PWT observation for that country (in 10 cases, the survey begins in the last PWT observation year). In these cases, I substitute the last available PWT observation for that country (and its corresponding date). In the case of three countries there are two surveys past the last PWT observation. In those cases, I drop one observation for each country. In sum, the PWT sample consists of 132 country x year observations, with 25 representating data before the corresponding DHS date (22 of these being within two years). The UN data extend to 2006, and hence can match all of the 135 survey x year observations of my DHS data.

Accounts Main Aggregates Database measures of GDP in constant US dollars. The UN growth measure is identical to PWT and the standard deviation of real incomes is only slightly higher. Neither result is surprising, as PWT mostly extrapolates international dataset measures of growth, while the UN database, despite being nominally at market exchange rates, makes ad hoc PPP adjustments to levels.[14]

The results reported above understate the true magnitude of the discrepancy between PWT and the DHS, as the DHS product level regressions understate the ratio of growth to standard deviation of real income by including product level cross-sectional variation in the denominator, as explained earlier in equation (6)′. To this end, Table IV below presents DHS random effects regressions run across all products, together, following equations (7), (9) and (10) presented earlier.[15] With the product level shocks purged from the country level standard deviation by the use of a panel, the average growth to standard deviation across all 26 products is now seen to be about .049, i.e. on the order of 2.5 times that present in PWT. This result is not due to the dominant influence of a particular product group as, with the exception of housing (brought down by the outlier of tap water in Table II), most product groups indicate a growth to standard deviation ratio about 2.5 to 3 times that of PWT. While the random effects at the product and product group level are generally significant and have some influence on the coefficient estimates and their standard errors, they are not essential to this result, as can be seen by comparing the different columns of the table.

[14] In the case of economies with volatile price levels and exchange rates, an adjustment is made using relative domestic/US inflation rates back to "the year closest to the year in question with a realistic GDP per capita US Dollar figure" (http://unstats.un.org/unsd/snaama/formulas.asp).

[15] To get things started, column (1) is actually equation (7) minus the country x product level random effect, i.e. the most basic panel formulation possible. This regression is inefficient, as the large country x product specific error is included in the general error term, but, relative to the results presented in Table III, accomplishes the basic task of averaging out, i.e. eliminating, the country x product specific shocks in the implicit calculation of the cross-sectional variation.

Table IV: DHS Estimates of Growth/Standard Deviation
Dependent variable = urban/rural weighted country means
Panel data: product x country x time observations
$y_{pct} = a_p + b_p*(gt + e_pt + e_Gt + u_c + e_{pc} + e_{Gc} + e_{pct})$

| | | (1) | (2) | (3) | (4) |
|---|---|---|---|---|---|
| All products | $g$ | .049 (.007) | .048 (.005) | .049 (.007) | .049 (.007) |
| | $\sigma_{pc}$ | | 1.20 (.124) | 1.19 (.123) | 1.13 (.122) |
| | $\sigma_p$ | | | .020 (.004) | .020 (.004) |
| | $\sigma_{Gc}$ | | | | .438 (.073) |
| | $\sigma_G$ | | | | .001 (.028) |
| | $\sigma_{pct}$ | 1.19 (.118) | .429 (.044) | .411 (.042) | .414 (.044) |
| Consumer durables | $g$ | .041 (.009) | .049 (.006) | .050 (.009) | |
| | $\sigma_{pc}$ | | .988 (.119) | .979 (.117) | |
| | $\sigma_p$ | | | .018 (.007) | |
| | $\sigma_{pct}$ | .900 (.095) | .304 (.036) | .291 (.034) | |
| Housing | $g$ | .022 (.008) | .026 (.004) | .026 (.008) | |
| | $\sigma_{pc}$ | | .841 (.107) | .848 (.107) | |
| | $\sigma_p$ | | | .015 (.006) | |
| | $\sigma_{pct}$ | .718 (.076) | .261 (.031) | .243 (.029) | |
| Children's Health | $g$ | .053 (.011) | .055 (.008) | .055 (.009) | |
| | $\sigma_{pc}$ | | .863 (.110) | .863 (.110) | |
| | $\sigma_p$ | | | .011 (.008) | |
| | $\sigma_{pct}$ | 1.02 (.113) | .588 (.070) | .582 (.070) | |
| Family Economics | $g$ | .062 (.014) | .064 (.010) | .062 (.013) | |
| | $\sigma_{pc}$ | | 1.83 (.279) | 1.80 (.273) | |
| | $\sigma_p$ | | | .023 (.009) | |
| | $\sigma_{pct}$ | 1.54 (.176) | .618 (.091) | .602 (.088) | |

Notes: All equations include a random effect at the country level ($u_c$) normalized to have a standard deviation of 1. Subscripts indicate the categories within which the shocks operate and across which the shocks are uncorrelated. The p and G random effects are growth shocks, i.e. multiplied by the time trend. The product group regressions cannot be executed with group level and growth random effects (Gc and G), as these would be colinear with the other terms in the regression. All equations include the first step covariance matrix as an additive part of the second step covariance matrix. See Section III for further details.

The results reported above use a particular functional form, the logit, to evaluate the dichotomous variables that form most of my product sample. The reader might be concerned that the ln odds ratio that this functional form produces transforms variation in mean values to variation in measured consumptions in such a way as to produce the results reported above. To explore this, I consider, as alternatives, the probit, weibull, cauchy and linear probability models. The predicted probabilities and estimated dependent variables associated with each functional form are presented in Table V. The probit has slightly thinner tails than the logit, the cauchy has dramatically thicker tails, the weibull is not even symmetric[16], and the linear model, of course, is simply a linear regression. The coefficient of variation, skewness and kurtosis of the 2912 product x country x survey consumption level estimates for dichotomous variables produced by these functional forms differ substantially. However, as shown in the bottom row of the Table, when the random effects regression of the upper right hand corner of Table IV earlier is estimated, the different functional forms yield virtually identical results. The relative time series and cross sectional variation present in the DHS consumption data is insensitive to the functional form used to convert means to consumption indices.

Overwhelmingly, that is across most products and product groups and across a number of potential functional forms, the ratio of growth to the standard deviation of real income, as measured by the real material consumptions recorded in the DHS, is dramatically higher than that present in the PWT. In the sections which follow I extend my methodology to allow the separate

[16] Symmetry being that $1 - F(y) = F(-y)$.

| Table V: Alternative Functional Forms | | | | | |
|---|---|---|---|---|---|
| | Logit | Probit | Weibull | Cauchy | Linear |
| Predicted Probabilities | $\frac{e^y}{1+e^y}$ | $\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{y} e^{-v^2/2}\text{dv}$ | $e^{-e^y}$ | $\frac{1}{\pi}\tan^{-1}(y)+\frac{1}{2}$ | N.A. |
| Dependent Variable | $\ln\left[\frac{\bar{x}}{1-\bar{x}}\right]$ | $F^{-1}(\bar{x})$ | $\ln(-\ln(\bar{x}))$ | $\tan[\pi(\bar{x}-1/2)]$ | $\bar{x}$ |
| Coefficient of Variation | 4.87 | 5.60 | 3.16 | 8.58 | 0.65 |
| Skewness | -0.57 | -0.38 | -0.47 | -18.23 | -0.08 |
| Kurtosis | 3.01 | 2.36 | 3.35 | 443.40 | 1.63 |
| $\hat{g}$ | .050<br>(.008) | .050<br>(.008) | .049<br>(.008) | .048<br>(.008) | .050<br>(.007) |

Notes: N.A. - not applicable, the linear model is not a predicted probability (between 0 and 1), but simply a linear regression. $F^{-1}$, inverse cumulative standard normal. Coefficient of Variation - average second central moment (of the 2912 estimated product x country x survey values for dichotomous variables) divided by the absolute value of the mean; Skewness & Kurtosis - average third and fourth central moments divided by the standard deviation raised to the third and fourth power, respectively. $\hat{g}$ = estimated value using specification of column (4) in Table IV across all products (including ln rooms, weight and height, estimated in all specifications as sample means).

estimation of both the growth rate and the standard deviation, thereby addressing the question of whether the DHS data suggest that the PWT growth rates are too low or the PWT standard deviation too high. As will be seen, my answer to this question focuses mostly on the growth rate, but individuals with strong priors concerning the return to education would produce different answers. It does not appear possible, however, to escape from the basic inconsistency highlighted in Tables II-V. All remaining arguments simply revolve around where, in the growth or standard deviation, one wishes to place the discrepancy.

## V. Methods II: Incorporating Micro Correlations

Let the real demand by household h for product p in region r in period t be described by the equation:

$$(11) \quad \ln(Q_{hprt}) \;=\; \alpha_p \;+\; \eta_p \ln(Y^R_{hrt}) \;+\; \vec{\beta}^{\prime}_p \vec{X}_{hrt} \;+\; \eta_p \varepsilon^{\overline{P}}_{prt} \;+\; \varepsilon_{hprt}$$

which is merely (4) earlier above estimated at the micro household level and augmented with a vector of household demographic variables $\vec{X}_{hrt}$ that shift demand through the coefficients $\vec{\beta}_p$. The error term is now made up of two components, the influence of relative prices, whose effect, as explained earlier, is proportional to the quasi-income elasticity of demand, and a mean zero idiosyncratic household preference shock. Ln household permanent real income per adult can reasonably be thought of as being given by:

$$(12) \quad \ln(Y^R_{hrt}) \;=\; \ln\left(Y^{R\sim E}_{rt}\right) \;+\; R_E E_{hrt}$$

where $E_{hrt}$ is the average years of educational attainment of adult household members, $R_E$ is the return to a year of education, and $\ln\left(Y^{R\sim E}_{rt}\right)$ is education adjusted ln regional real income at time t. Ln average regional real income at time t is given by $\ln(Y^R_{rt}) \;=\; \ln(Y^{R\sim E}_{rt}) + \; R_E \overline{E}_{rt}$, where $\overline{E}_{rt}$ is mean educational attainment, and, as before, is assumed to be growing at a common global growth rate, $\ln(Y^R_{rt}) \;=\; \ln(Y^R_r) \;+\; gt$.

For each product, combine a number of household level surveys to estimate the equation

$$(13) \quad \ln(Q_{hprt}) \;=\; \hat{a}_{prt} \;+\; \hat{b}_p R_E E_{hrt} \;+\; \vec{\hat{c}}^{\prime}_p X_{hrt} \;+\; e_{hprt}$$

where the $\hat{a}$s are a complete set of product specific region x time (or, equivalently, survey) dummies. Regions, r, can be defined at any level that allows consistent aggregation across time, and, in my case, will consist of the urban and rural areas of each country. Clearly, the estimates $\hat{b}$ and

$\vec{c}$, identified off of cross sectional variation within surveys, will be unbiased, but the dummies will capture all common product x region x time components:

$$(14) \quad \hat{b}_p = \eta_p \qquad \hat{\vec{c}}_p = \vec{\beta}_p \qquad \hat{a}_{prt} = \alpha_p + \eta_p \ln(Y_{rt}^{R\sim E}) + \eta_p \varepsilon_{prt}^{\vec{P}}$$

While the unconditional expectation of $\varepsilon_{prt}^{\vec{P}}$, the influence of relative prices, is zero, it takes on particular values within any particular product x region x time grouping and ends up being incorporated into the constant term.

Finally, construct measures of real regional income at time t, as implied by consumption of a particular product, as the sum of the product x region x time first stage dummy divided by the income elasticity of demand, plus the ln real income attributable to the separately estimated average regional educational attainment:

$$(15) \quad \ln(\hat{Y}_{prt}) = \frac{\hat{a}_{prt}}{\hat{b}_p} + R_E \hat{\bar{E}}_{rt}$$

When weighted and added using regional (i.e. urban/rural) weights, these measures produce a panel dataset of country real incomes, as measured by different product equations. These measures can then be projected, in a random effects panel regression on product dummies, growth rates, and random effects designed to capture cross correlations in the error term brought about by levels and trends in relative prices:

$$(16) \quad \ln(\hat{Y}_{pct}) = \hat{a}_p + \hat{g}t + u_c + e_p t + e_G t + e_{pc} + e_{Gc} + e_{pct}$$

The reader will recognize this as no more than equation (10) earlier in Section III divided by the quasi-income elasticity of demand, which is now estimated in the first stage regressions which product the estimates of the dependent variable. Asymptotically

$$(17) \quad \hat{a}_p = \alpha_p/\eta_p \quad , \quad \hat{g} = g \quad , \quad \hat{\sigma}_u = \sigma_{Y^R} .$$

where $\sigma_{Y^R}$ is the standard deviation of country real incomes, measured as the urban/rural weighted average of permanent regional incomes ($\ln(Y_r^R)$ earlier). It should be apparent that (16) can be run using the original urban or rural income estimates from (15), the same estimates without mean educational adjustment (i.e. simply $\hat{a}_{prt}/\hat{b}_p$), or even the difference of such estimates, to produce estimates of the growth and standard deviation, with or without adjustment for educational attainment, of urban and rural incomes and the urban-rural difference. These can then provide insight into the factors behind the growth and overall cross-country dispersion in real incomes, as estimated by the country level version of (16).

## (a) Details and Extensions

A quick examination of equations (13) and (15) above reveals the obvious fact that the quasi income elasticity, $\hat{b}_p$, is inversely related to the return to education, $R_E$, so that $\ln(\hat{Y}_{prt})$, real regional x time consumption income as defined by product p, is linear in $R_E$. It follows that all of the equations, up through and including the random effects panel regression, (16), used to estimate the growth and standard deviation of income, can be estimated in terms of years of education (i.e. setting $R_E$ temporarily equal to 1) and the results then multiplied by one's estimate of $R_E$ to arrive at income equivalent measures of growth or variation. I highlight this fact so that the reader can see that any disagreement with my estimate of $R_E$ from the DHS data (later below) can be resolved by simply scaling the estimates and standard errors proportionately to the preferred number.

As in Section III, in estimating (16) I add the first step covariance matrix for the dependent variable, $\ln(\hat{Y}_{pct})$, to the random effects covariance matrix in calculating the covariance matrix for the GLS likelihood. This is not only justifiable econometrically, as a means of improving the efficiency of the estimates, but also, in more pedestrian terms, makes the procedure somewhat "idiot proof". Products where the estimated relationship between relative incomes (education)

and demand is statistically weak will have very large first step standard errors. In this regard, putting aside the random effects covariance terms, the GLS likelihood will function like weighted least squares, discounting the variation in those observations. Consequently, randomly picked real choice variables, such as the household's favourite colour, whose association with real incomes is dubious, will have little effect on the estimates of real relative incomes.[17]

The inclusion of the first-step covariance matrix for $\ln(\hat{Y}_{pct})$ in the second step GLS likelihood raises an important technical stumbling block. As shown in (15) above, $\ln(\hat{Y}_{prt})$ is computed as the ratio of normally distributed variables. In calculating the distribution of non-linear functions of normal variables, it is customary to make use of the "delta method", an application of the central limit theorem. However, even the central limit theorem has its limits. As the probability mass around zero of the random variable in the denominator increases, the central limit theorem breaks down, the most notable example of which is the well known result that the ratio of two independent standard normal variables follows a cauchy distribution, which doesn't even have any moments. To the degree that the denominator in (15), the quasi-income elasticity, differs from zero, this is not a problem as, asymptotically, the probability mass of the estimated coefficients around zero goes to zero. However, for finite samples, or in the case of poorly chosen variables whose correlation with real incomes is spurious, the probability mass around zero can be large enough to make the delta method calculation of the first step covariance matrix utterly inaccurate. I handle this problem by using Monte Carlo techniques to estimate the first

[17] However, the estimates are not protected from bias introduced by the use of real variables, such as race, which are strongly correlated with education and incomes but not subject to individual choice.

step covariance matrix of $\ln(\hat{Y}_{prt})$.[18] In most cases, and for all the baseline estimates, this produces first step covariance matrices which match, virtually identically, those calculated using the delta method. However, for more esoteric extensions, with smaller estimated quasi-income elasticities and larger first step standard errors, the Monte Carlo techniques occasionally produce standard errors for $\ln(\hat{Y}_{prt})$ which are 1000s of times larger than those estimated using the delta method.[19] This affects the results and, in what follows, is reported where relevant.

Household survey data are collected in "clusters", i.e. groups of households at particular survey locations. This suggests the likelihood of correlation across the error terms for households in the same cluster which makes the first step covariance matrix inaccurate and the coeffi-

[18] To be clear, I accept the standard maximum likelihood estimates of the first step coefficients and their covariances, as these are based upon 100,000s of observations and do not involve ratios of normals. But, in calculating the distribution of (15), which involves the ratios of these normal variables, I generate a million draws from the estimated joint distribution of the multiple $\hat{a}_{prt}$'s and single $\hat{b}_p$ in each product equation and then calculate the resulting distribution of the ratios $\hat{a}_{prt}/\hat{b}_p$, to which I then add the covariance matrix of $R_E$ times the estimated mean educational attainment by region.

[19] The reader will no double have immediately noted that, in these cases, the use of the Monte Carlo estimates in the second step is not exactly valid. When the Monte Carlo covariance matrix closely matches that computed by the delta method, the distribution can be taken as normal and either covariance matrix used in the second step. However, when the Monte Carlo covariance matrix differs substantially from the delta method, the resulting distribution cannot be taken as normal and mechanically pushed into the normal covariance matrix of the second step calculations. I make this adjustment on the basis of the conjecture that a correction for extremely large second moments, while not taking into account the changes to higher moments, is better than none at all. I intend to confirm this conjecture with simulations. In statistical theory, there has been little progress on distributions of this type, beyond extending the Cauchy result to noting that the distribution of the ratio of two correlated normals produces a distribution that takes close to a page of text to write down. Thus, analytically calculating the finite sample multivariate distribution of (15) is not an option.

I should note that this problem does not invalidate the statement earlier above that the incorporation of the first step covariance matrix in the second step likelihood protects the estimates from "idiot" variables which are not actually correlated with real incomes. As the probability mass of the estimated quasi-income elasticity around zero grows, the estimated variance explodes to infinity; it simply explodes incredibly faster than indicated by the central limit approximation, as I find in my Monte Carlo simulations.

cient estimates inefficient or, worse, biased. I address this problem in three ways. First, as a baseline, I ignore the clustering in the first step estimation procedure, but calculate more accurate "clustered" first step standard errors using the usual sandwich covariance estimator. Second, I formally estimate first step cluster-level random effects regression and discrete choice logit models. Third, to allow for the possibility that the cluster errors are correlated with the independent variables, I estimate cluster-level fixed effects models.[20]

Each successive cluster level model I apply is, statistically speaking, found to be superior to the one before, i.e. no random effects are rejected in favour of significant random effects and random effects (and their assumption of independence from the other regressors) are rejected in favour of fixed effects. The large correlation of errors within clusters in the random effects specification places greater weight on "within" cluster variation in educational attainment and consumption levels, which the fixed effects specification completes by looking only within clusters. This produces, empirically, smaller estimates of the quasi-income elasticity of demand and, by extension, greater estimates of the cross sectional and time series variation in living standards. However, it is not clear these estimates are an improvement on those found ignoring cluster level correlations. First, as one tunnels down to the cluster level, the noise to signal ratio in measures

[20] For the dichotomous variables, I use Butler & Moffitt's (1982) random effects specification, modelling the random effect as normally distributed and using Gauss-Hermite quadrature to integrate the cluster joint logit probability and, for fixed effects, Chamberlain's (1980) conditional logit likelihood, implicitly differencing out the cluster fixed effect (without estimating it) by evaluating the likelihood of a particular cluster outcome conditional on overall cluster characteristics.

As, for both logit and regression, the regional dummies cannot be directly estimated with cluster fixed effects, I employ a two step procedure: first, estimating the income elasticity and demographic coefficients using cluster fixed effects, and then using these estimated coefficients as an offset in a cluster random effects specification where I calculate the regional product dummies. This is justified on the obvious grounds that the cluster errors, within regions, are orthogonal to the regional means. The standard errors of the regional dummies and their covariance with the estimated income elasticity are adjusted for the two-step procedure.

of household educational attainment rises, biasing the coefficients towards zero. Thus, it is not clear whether the smaller estimates of quasi-income elasticities of demand are more accurate representations of reality. Second, much of the correlation within clusters in consumption represents, in fact, the outcome of demand (for communal infrastructure) that is implicitly paid for through the cost of housing and land. To this extent, one would clearly want to identify the quasi-elasticity of demand using between cluster, rather than within cluster, variation. For these reasons, I treat estimates without adjustment for cluster random or fixed effects as my baseline,[21] reporting the others as variations for the reader's consideration.

It is obvious that there are likely to be strong regional variations in demand patterns or the relationship between income, consumption and outcomes. Two obvious examples are the demand for tap water (depending on its quality) and the relationship between nutrition and childhood weight and height (depending on genetics). To the degree that these affect the level of consumption, independent of income, their undue influence on the estimates is eliminated by the product and product group level random effects, which evaluate the degree to which consumption levels are correlated within product and product groups in a country, but not across all products in that country.[22] However, an alternative problem is the possibility of local level

[21] In this case, the correlation within clusters influences the standard errors, which recognize that it diminishes the effective size of the sample, but is not allowed to influence the coefficient estimates.

[22] Obviously, these random effects do not eliminate the influence of such factors, they simply change the way the product and product group observations are weighted. However, as explained earlier, larger random effects of this type imply a smaller overall estimate of the country level random effect (the cross sectional variation), as the estimates correct for the fact that more of the observed country level variation is simply due to the small number of independent product/product group level observations.

variation in the estimated slopes, i.e. in the relation between incomes and consumption. To this end, consider reformulating demand as being given by:

$$(18) \quad \ln(Q_{hprt}) = \alpha_p l^{\alpha}_{prt} + \eta_p l^{\eta}_{prt} \ln(Y^{R}_{hrt}) + \vec{\beta}'_p l^{\beta}_{prt} X_{hrt} + \eta_p l^{\eta}_{prt} \varepsilon^{\vec{P}}_{prt} + \varepsilon_{hprt}$$

where $l^{i}_{prt}$ represents an unmeasured regional factor influencing the demand coefficients i. Next consider estimating each product demand equation survey x region by survey x region and using the survey x region-specific income elasticities of demand to calculate the real regional income implied by the consumption of a particular product

$$(19) \quad \ln(Q_{hprt}) = \hat{a}_{prt} + \hat{b}_{prt} R_E E_{hrt} + \hat{\vec{c}}'_{prt} X_{hrt} + e_{hprt}$$

$$(20) \quad \hat{Y}_{prt} = \frac{\hat{a}_{prt}}{\hat{b}_{prt}} + R_E \hat{\bar{E}}_{rt} = \frac{\alpha_p l^{\alpha}_{prt}}{\eta_p l^{\eta}_{prt}} + \varepsilon^{\vec{P}}_{prt} + \ln(Y^{R\sim E}_{rt}) + R_E \hat{\bar{E}}_{rt}$$

Clearly, this equation merely adds additional error (stemming from the level variation in consumption patterns) to the product level measures of real regional consumption, so that estimating equation (16) using (20) as data asymptotically yields measures of the growth and standard deviation of real incomes. I show, later in the paper, that empirically this produces estimates that reinforce, rather than reverse, the main results of the paper, with somewhat smaller point estimates of the standard deviation of income and somewhat larger estimates of its growth rate, although the standard errors are quite large. There is definitely statistically significant variation in demand patterns across regions, but it is not very systematically correlated with the variables of interest, so that expansion of the procedure to calculate idiosyncratic regional demand patterns has a large effect on standard errors and a comparatively small effect on coefficient estimates.[23]

[23] This is the "esoteric" variation alluded to above, where Monte Carlo simulations produce dramatically larger standard errors for the ratios of first step coefficients. This stems from the fact that, with the quasi income elasticities estimated region by region, there are a number of cases where the t-statistics are rather low (i.e. below 3) and the cauchy-like distribution comes into play.

## VI. Results II: The Standard Deviation & Growth Rate of Living Standards

### (a) The Return to Human Capital

As a preliminary, I use DHS data on individual earnings from work to calculate the return to education. I focus on individuals 25 or older, whose education can be taken as completed, reporting earnings from working for others (i.e. not for family or self). I find earnings data of this sort for adult women in 26 DHS surveys in 14 sub-Saharan African and 10 other countries, and for adult men in a sub-sample of 16 of these surveys in 11 sub-Saharan countries and 5 other countries (see Appendix I). I run the typical Mincerian regression of ln wages on educational attainment, age, sex and regional controls.

As shown in Table VI, the OLS estimate of the return to human capital is somewhat sensitive to the number and level of regional controls. While column (1) includes the most basic controls, a dummy variable for the nominal level of wages in each survey, column (2) includes survey x rural/urban controls. Doubling the number of locational controls in this fashion lowers the return to a year of education[24] from 11.5 to 10.8 percent. Adding random effects at the cluster level lowers the marginal return further, while fixed effects at the cluster level (column 4) bring it down to 9.5 percent. These results can be rationalized by arguing that rich people tend to live together in rich places, i.e. regions and locales (such as urban centres) which provide higher earnings for any given level of education. As more detailed locational controls are introduced, the return to education is increasingly identified from within locale differences in educational attainment and incomes, rather than cross regional income differences. However, it is also important to note that more detailed locational controls increase the noise to signal ratio in educational attainment, biasing the coefficient towards zero. This is particularly relevant for the

[24] As noted in appendix I, I use the DHS data to construct measures of educational attainment, not simply years of attendance.

| Table VI: Ln Wage Regressions | | | | | |
|---|---|---|---|---|---|
| | (1) survey dummies | (2) survey x rural/urban dummies | (3) cluster random effects | (4) cluster fixed effects | (5) cluster fixed effects (IV) |
| educ | 11.47 (.154) | 10.78 (.148) | 10.40 (.127) | 9.46 (.156) | 11.60 (.477) |
| age | 4.73 (.692) | 4.68 (.675) | 4.93 (.614) | 4.81 (.701) | 4.58 (.819) |
| $age^2$ | -0.05 (.009) | -0.05 (.009) | -0.05 (.008) | -0.05 (.009) | -0.04 (.011) |
| sex | -35.04 (1.95) | -35.95 (1.91) | -36.52 (1.52) | -36.65 (1.70) | -39.56 (1.99) |
| N | 22996 | 22996 | 22996 | 22996 | 18418 |

Notes: Dependent variable = 100*ln annualized wage income of individuals 25 or older working for others, so the coefficients can be read as derivatives expressed in percent. Educational attainment and age measured in years, while sex = 1 if female.

estimates with cluster fixed effects, as these dummies account for 58 percent of the residual (orthogonal to the individual controls) variation in individual educational attainment.

Column (5) of Table VI controls for the measurement error in individual educational attainment by instrumenting it with the mean educational attainment of other adult members of the same household, as well as their mean age, age2 and sex.[25] As shown, when instrumented, the estimated return on human capital jumps to 11.6 percent. When compared with the coefficient for column (4), this suggests that measurement error accounts for about .19 of the residual variation in individual educational attainment in that specification.[26] This would imply a measurement standard error of about 1.6, i.e. that about 36 percent of the wage reporting sample,

[25] The absolute values of the t-statistics of these four variables in the first stage regression are 90.5, 3.4, 6.2, and 13.4, respectively.

[26] As is well known, attentuation bias when one variable is measured with error is equal to S/(S+N), where S is the orthogonal (to other regressors) signal variation and N is the noise variation. In the sentence which follows, I multiply 1 minus this ratio times the residual variance estimate (adjusted for degrees of freedom) of the regression of E on cluster fixed effects and the other individual level controls. I should note that the sample of column (5) is smaller, because many individuals live in households without other adults, and hence cannot be instrumented. The coefficient of column (4) using the sample of column (5) (which is what I use to calculate the noise to signal ratio) is 9.40.

with mean educational attainment of 9.5 years,[27] over or understate their educational attainment by 1.6 years or more. This is large, but by no means implausible. Adjusting the coefficient of column (2) by this estimate of measurement error produces a point estimate of the return to education of 12.7 percent in that column.[28] When compared with column (5)'s point estimate, this indicates that while measurement error is a concern, there is also substantial correlation, below the urban rural level, between individual's incomes and the education-adjusted income level of the locales they live in.

In what follows, I will take 11.6 percent as my "known" estimate of $R_E$. Psacharopolous (1994) in his oft cited survey of Mincerian regressions, finds an average marginal return of 13.4 percent in 7 studies of sub-Saharan Africa and 12.4 in 19 studies of Latin America and the Caribbean, regions which, together, make up 3/4 of the countries in my sample. Thus, the number I use is not particularly large or out of keeping with the existing literature. However, readers who have strong alternative priors can simply scale all of the growth rates and standard deviations of real incomes reported below by the ratio of their preferred number to 11.6. This does not change the estimated ratio of growth to standard deviation, but does change one's evaluation of where the discrepancy between the DHS and international data sets, such as PWT, lies.

[27] The wage reporting sample is considerably better educated than the average for the men and women in the male & female survey modules from which the wage data come (5.3 years). Although, this suggests the possibility of selectivity bias, when I implement the standard selectivity bias adjustments I do not find that this substantially changes the estimated return to human capital. However, I still need to explore this further.

[28] Arrived at by calculating the estimated residual variation of E when regressed on the individual controls and column (2)'s fixed effects and subtracting the measurement error variance noted above.

## (b) First Step Estimates

Table VII below reports the coefficients on household mean years of adult educational attainment in product by product demand equations, estimated with country x survey x urban/rural dummies and household and individual demographic controls, as listed in the notes to the table. With the exception of ln weight, height and rooms per capita, the dependent variable in each row is a 0/1 dichotomous variable and the reported figures represent the coefficients in a logit discrete choice model. The second and third columns of the table run the baseline specification with cluster random and fixed effects, which, as noted earlier, tend to lower, somewhat, the absolute value of the education coefficient. The fourth column runs a specification, described later in section VIII, with household level random effects. In this case all of the equations are linked through the random effect and estimated simultaneously. This has a substantial influence on the education coefficients, raising the absolute value of some and lowering others. In regards to producing estimates of the levels of real living standards, however, these changes average out, as will be seen later in the paper.

For our purposes, the main relevance of Table VII is that it establishes that each of the real consumption variables used in this paper is very significantly and, generally, quite substantially related to real income, as measured by years of education. Across the different specifications, only one coefficient (marriage of young women with cluster fixed effects) is even close to being insignificant at the 1% level. The baseline income elasticities, evaluated at the sample

Table VII: Product Level Estimates of the Response to Educational Attainment
(= $\hat{b}_p * R_E$ in (13) above)

| | (1) baseline | (2) cluster random effects | (3) cluster fixed effects | (4) household random effects | (5) baseline Y elasticity |
|---|---|---|---|---|---|
| Radio | .153 (.001) | .149 (.001) | .134 (.001) | .160 (.001) | 0.56 |
| Television | .236 (.001) | .220 (.001) | .192 (.001) | .407 (.002) | 1.21 |
| Refrigerator | .253 (.001) | .236 (.001) | .202 (.001) | .441 (.002) | 1.64 |
| Bicycle | .056 (.001) | .077 (.001) | .078 (.001) | .059 (.001) | 0.34 |
| Motorcycle | .190 (.001) | .200 (.001) | .193 (.001) | .188 (.001) | 1.47 |
| Car | .250 (.001) | .234 (.001) | .191 (.001) | .282 (.002) | 2.01 |
| Telephone | .248 (.001) | .227 (.001) | .192 (.001) | .396 (.002) | 1.77 |
| Electricity | .228 (.001) | .235 (.002) | .216 (.001) | .375 (.002) | 0.93 |
| Tap Drinking Water | .076 (.001) | .057 (.001) | .046 (.001) | .093 (.001) | 0.36 |
| Flush Toilet | .234 (.001) | .224 (.002) | .196 (.001) | .266 (.002) | 1.37 |
| Constructed Floor | .210 (.001) | .207 (.001) | .185 (.001) | .253 (.002) | 0.73 |
| ln(Rooms/Capita) | .020 (.000) | .015 (.000) | .012 (.000) | .023 (.000) | 0.17 |
| ln Weight | .007 (.000) | .006 (.000) | .005 (.000) | .007 (.000) | 0.06 |
| ln Height | .002 (.000) | .002 (.000) | .001 (.000) | .003 (.000) | 0.02 |
| No Diarrhea | .033 (.001) | .032 (.001) | .021 (.001) | .030 (.001) | 0.06 |
| No Fever | .019 (.001) | .019 (.001) | .014 (.001) | .022 (.001) | 0.05 |
| No Cough | .006 (.001) | .008 (.001) | .005 (.001) | .015 (.001) | 0.02 |
| Alive | .059 (.002) | .059 (.002) | .046 (.002) | .034 (.001) | 0.04 |
| At School (6-14) | .200 (.001) | .171 (.001) | .151 (.001) | .168 (.001) | 0.50 |
| At School (15-24) | .148 (.001) | .135 (.001) | .111 (.001) | .135 (.001) | 0.89 |
| Working (15-24) | -.032 (.002) | -.037 (.002) | -.042 (.003) | -.022 (.001) | -0.16 |
| Working (25-49) | .020 (.001) | .028 (.001) | .025 (.001) | .006 (.001) | 0.08 |
| Birth (15-24) | -.012 (.001) | -.012 (.001) | -.007 (.002) | -.013 (.001) | -0.07 |
| Birth (25-49) | -.026 (.001) | -.024 (.001) | -.011 (.001) | -.033 (.001) | -0.19 |
| Marriage (15-24) | -.058 (.001) | -.035 (.001) | -.002 (.001) | -.100 (.001) | -0.28 |
| Marriage (25-49) | -.077 (.001) | -.064 (.001) | -.025 (.002) | -.069 (.001) | -0.04 |

Note: The reported number is the coefficient (standard error) on household mean adult educational attainment in years, with each equation including a complete set of country x survey x region (urban/rural) dummies and the following controls: (1) consumer durables & housing: ln number of persons in the household; (2) children's health: sex, ln(1+age in months) and ln(1+age in months) squared (for all but height and weight, which are quite linear in ln(1+age)); (3) household economics: age and age squared, as well as sex for education attendance variables (all others refer to women alone). Each equation is estimated separately, except for column (4), where the household level random effect (measured in equivalent years of education) links them all (see Section VIII further below).

mean probability,[29] coupled with the standard deviation (s.d.) of mean household adult education (4.5 years) and implied s.d. of predicted incomes (4.5*.116 ≅ .5) produces substantial variation in predicted outcomes.  Thus, a one s.d. movement in educational attainment produces a ln 28 percent higher relative probability of owning a radio (mean value of .574 - see Table I) and a ln 69 percent higher probabily of having a flush toilet (.322).  Given the early age of the subjects (0-35 months), children's weight and height move relatively less, an average of 3 and 1 percent, respectively, with a s.d. movement in educational attainment, but are, nevertheless, very significantly correlated with household incomes.  The cumulative probability of survival for the average 0 to 35 month year old (mean value of .930) rises 2 percent with a s.d. movement in predicted incomes, a small apparent movement, but actually an implied fall in average cumulative mortality from .07 to .05.  The probability children and youths are in school rises 25 percent (mean value of .712) and 45 percent (.304) with a s.d. movement in incomes, while the probability a young woman is working (.416) or ever-married (.431) falls by 8 percent and 14 percent, respectively.  The income elasticities implied by the coefficients in the other columns, in some cases higher and in some cases lower, can be arrived at by multiplying column (5) by the ratio of each column's coefficient to that listed in column (1).

[29] For the ln variables (weight, height and sleeping rooms), the implied income elasticity is B/.116, where B is the coefficient.  For the logit dichotomous variables, the elasticity of the probability with respect to real income is B*(1-P)/.116, where P is the mean value (Table I).

### (c) The Growth and Standard Deviation of Real Consumption

Table VIII below presents second step estimates of the growth and standard deviation of living standards using the baseline product level estimates of income elasticities and product x region x time constant terms described above in Section V to produce the dependent variable. In the top panel of the table, successive random effects are added, controlling for correlation at the product and product group level within countries and in growth rates. These adjustments tend to systematically lower the estimated cross sectional variation in incomes, while moving the estimated growth rate ever so slightly up and down. They are clearly, however, not crucial. Overall, the DHS data suggest a level of cross sectional variation consistent with that present in PWT measures of consumption per capita or per equivalent adult (between .64 and .67, in Table III earlier), but the estimated DHS growth rate is on the order of 2.5 to 3 times as large as the PWT growth rate (.011 to .013). The bottom panel of the table calculates the same measures at the product group level, showing that this basic pattern is present in virtually all product group categories. The slowest growing categories, housing and family economics, all indicate real consumption growth on the order of double that present in the PWT, and, aside from housing, all product groups show a level of cross sectional variation roughly equivalent to that present in the PWT.

Table IX explores the sensitivity of the results to various econometric techniques and functional form assumptions. First, in the upper left panel, I present results where the covariance matrix of the estimated dependent variables is not incorporated in the second step GLS likelihood. As shown, this dramatically raises both the growth and estimated standard deviation as the procedure no longer corrects for the fact that much of the cross sectional and time series variation in the dependent variable comes from the error in the first step estimates, particularly in the estimate of the income elasticity which produces correlated expansions and contractions of

Table VIII: DHS Estimates of the Growth and Standard Deviation of Living Standards
Dependent variable = urban/rural weighted country means

All products combined
$y_{pct} = a_p + gt + e_p t + e_G t + u_c + e_{pc} + e_{Gc} + e_{pct}$

| | | | | |
|---|---|---|---|---|
| $g$ | .030 (.004) | .032 (.001) | .030 (.004) | .030 (.005) |
| $\sigma_u$ | .741 (.073) | .706 (.071) | .705 (.071) | .698 (.073) |
| $\sigma_{pc}$ | | .866 (.020) | .871 (.020) | .834 (.020) |
| $\sigma_p$ | | | .019 (.003) | .018 (.003) |
| $\sigma_{Gc}$ | | | | .278 (.042) |
| $\sigma_G$ | | | | .006 (.005) |
| $\sigma_{pct}$ | .887 (.014) | .268 (.006) | .252 (.006) | .252 (.006) |

By product group
$y_{pct} = a_p + gt + e_p t + e_G t + u_c + e_{pct}$

| | Consumer Durables | Housing | Children's Health | Family Economics |
|---|---|---|---|---|
| $g$ | .046 (.010) | .021 (.008) | .031 (.004) | .024 (.004) |
| $\sigma_u$ | .735 (.089) | 1.07 (.123) | .574 (.068) | .592 (.071) |
| $\sigma_{pc}$ | .969 (.042) | 1.01 (.054) | .503 (.030) | .763 (.036) |
| $\sigma_p$ | .024 (.007) | .016 (.006) | .005 (.005) | .009 (.005) |
| $\sigma_{pct}$ | .238 (.009) | .294 (.015) | .278 (.017) | .218 (.009) |

Table IX: Sensitivity Tests
$y_{pct} = a_p + gt + e_p t + e_G t + u_c + e_{pc} + e_{Gc} + e_{pct}$

first step logit for dichotomous variables

| | 2nd step w/out 1st step covariance | 1st step cluster random effects | 1st step cluster fixed effects | 1st step local coefficients |
|---|---|---|---|---|
| $g$ | 058 (.016) | .035 (.005) | .035 (.007) | .049 (.027) |
| $\sigma_u$ | .932 (.120) | .821 (.087) | .803 (.087) | .548 (.103) |

alternative first step functional forms

| | Probit | Weibull | Cauchy | Linear |
|---|---|---|---|---|
| $g$ | .029 (.004) | .028 (.004) | .039 (.008) | .027 (.003) |
| $\sigma_u$ | .666 (.070) | .665 (.070) | .958 (.106) | .644 (.068) |

Note: Each specification includes the full set of error terms (p, G, pc, Gc, pct, etc), as in the upper right panel of Table VIII, but only g and $\sigma_u$ are reported.

the dependent variables.[30] Turning to the second and third columns of the upper panel, these incorporate cluster level random and fixed effects in the first step equations used to produce the dependent variable. With somewhat smaller estimated income elasticities, on average, they expand the overall variation in the sample, producing estimates of the growth rate and standard deviation about 1/6 larger than those arrived at in the baseline calculations of Table VIII. The last column of the upper panel of Table IX makes use of local coefficient estimates of the first step parameters, including the income elasticity, in the manner described earlier in Section V. This produces greater growth and less cross sectional variation than the baseline estimates, but the standard errors are also substantially increased as a great deal of imprecision is added when each demand equation is estimated region by region. Finally, the bottom panel of the table uses different functional forms (as described earlier in Section IV) in the calculation of the first step estimates. The results are remarkably similar, with the exception of the cauchy which increases the relative change associated with differences and movements in the tails of the distribution, producing a 1/3 increase in the estimated growth rate and standard deviation of consumption.[31]

Table X recalculates the panel data model for the PWT and DHS, incorporating a dummy for the relative level and relative growth rate of sub-Saharan Africa. This table highlights further inconsistencies between the two data sources. Although the standard errors of both groups of

---

[30] This effect is even greater for the remaining random effects, fixed effects and cluster effects estimates reported in the table (which, as reported, incorporate the first step covariance matrix into the GLS likelihood) as the standard errors of their first step estimates are considerably larger.

[31] Relative to the logit, probit & weibull, the cauchy has dramatically thicker tails. Hence, any given difference or movement in the mean value of a random variable in the tails (e.g. a change in the mean ownership of cars from .05 to .1) is associated with a much greater movement in the index (relative to movements around a mean value of .5).

| Table X: The Relative Poverty of Sub-Saharan Africa | | | |
|---|---|---|---|
| | PWT 6.2 consumption per capita | equiv. adult | DHS baseline |
| *g* | .019 (.003) | .017 (.003) | .031 (.005) |
| Africa's relative growth | -.011 (.005) | -.009 (.005) | -.002 (.003) |
| Africa's relative level | -.732 (.153) | -.674 (.149) | -.972 (.147) |
| $\sigma_u$ | .566 (.054) | .553 (.053) | .502 (.057) |

Note: With A a dummy for Sub-Saharan Africa, PWT uses $y_{ct} = a + A + g*t + g_A*A*t + u_c + e_{ct}$, while DHS uses $y_{pct} = a_p + A + g*t + + g_A*A*t + e_pt + e_Gt + u_c + e_{pc} + e_{Gc} + e_{pct}$. For brevity, not all error terms and random effects are reported in the table.

estimates are quite large, the DHS suggest that Africa is substantially poorer than indicated by PWT, with a relative disadvantage in levels between 1/3 and 1/2 greater than that shown by PWT. While the PWT indicates that Africa is growing about 1 percent slower than the rest of the sample, for an average growth rate of .7 or .8 percent, the DHS indicates that there is virtually no difference in Africa's trend, which appears to be around 2.9 percent per annum! The relative gap between the non-African growth rates of the two samples (1.7 to 1.9 percent versus 3.1 percent), while still large, is comparatively smaller. Once the greater gap in sub-Saharan and non sub-Saharan levels is taken into account, the residual cross sectional variation in the DHS is about 1/10 smaller than that present in the PWT.

The results of this section can be summarized as follows: in the aggregate, most of the discrepancy between the DHS and PWT lies in the growth rate, where the DHS indicate 3 percent average growth of real consumption, as compared to the 1.1 to 1.3 percent implied by PWT. The cross sectional standard deviation of real consumpionn in the two sources, at about .65 to .7, is comparable. This greater inconsistency in growth, and relative similarity in cross sectional variation, is, broadly speaking, present across all product sub-groups within the DHS data. Alternative functional forms or econometric specifications for the DHS data generally suggest similar growth rates, with more esoteric functional forms or econometric specifications indicat-

ing substantially larger growth rates, increasing the discrepancy with PWT. A more detailed breakdown of the sample into sub-Saharan and non sub-Saharan countries indicates that the greatest relative growth discrepancy lies in sub-Saharan Africa, whose DHS growth rate is perhaps 3.5 times that indicated by PWT. The remaining countries have a growth rate about 1.7 times greater than indicated by PWT. While the PWT indicate that Africa's growth rate is considerably slower than that of the rest of the sample, the DHS do not. However, the DHS data indicate that sub-Sahara's relative income level is much worse than suggested by PWT. Once this is taken into account, residual cross sectional variation is somewhat lower in the DHS.

In the sections which follow, I show that part of this varied pattern of differences and discrepancies actually stems from a fundamental difference in the concept measured by the PWT and my DHS estimates, revolving around the notion of income inequality within regions and countries. In so doing, I show how the DHS data produce a coherent data set of the level, growth and local inequality of real consumption which, with the notable exception of African growth, is broadly consistent with the information in PWT and other international data sets.

## VII. Income Inequality, the Ln of the Mean and the Mean of the Ln

Conventional measures of real living standards calculate the average consumption per capita, i.e. total aggregate real consumption divided by the number of persons or equivalent adults. When the growth and cross-national dispersion of these measures is examined, the ln is usually taken, so they may reasonably be termed "the ln of the mean". In my use of the DHS data to produce estimates of real living standards, I project a measure of real consumption on years of educational attainment to derive quasi-income elasticities and implied real income values for the product x survey x urban/rural regional dummies. This procedure is justified by assuming that the demand measure is, at least approximately, linear in the ln of income, which is in turn linear in years of education. If so, then the estimated regional dummies are estimates of the average of education adjusted ln income. I then add back in average regional educational attainment to produce a product specific estimate of the average of ln real income (eqn. (15) earlier). Weighting these urban/rural estimates with the relative number of urban/rural households, I produce a product level measure of the national average ln income, which is then used in equations such as (16). In sum, my procedure can reasonably be described as estimating the "mean of the ln".

From Jensen’s inequality, we all know that the mean of the ln is less than the ln of the mean. What is more relevant, however, is that the gap is related to the degree of dispersion of the random variable; in our case, the dispersion of real incomes. To make things concrete consider, as a base, the case where ln household incomes, $\ln(Y^R)$, are ln normally distributed with mean $\mu$ and standard deviation $\sigma$. In this case we have the well known result:

$$(21) \quad E[\ln(Y^R)] = \mu \qquad \ln[E(Y^R)] = \mu + .5\sigma^2$$

Thus, conventional measures of living standards, such as PWT, are actually a mixture of the average of ln real living standards and the dispersion of the same.

In the sections which follow, I extend my methodology to allow the separate estimation of

both the average level and within region dispersion of ln household incomes.  This allows me to produce "ln of the mean" equivalent estimates of the levels and growth of real incomes.  As will be seen, these adjustments move my figures systematically closer to those of PWT, as would be expected from the argument presented above.  Beyond this, the approach allows me to produce an integrated data set of means and dispersions of regional incomes and education levels.

## VIII. Methods III: Estimating Income Inequality

As a final extension of my methodology, I begin by specifying, as before, that real demand by household h for product p in region r at time t is a function of real income, demographic variables, a common shock reflecting the impact of time specific regional relative prices, and idiosyncratic household level variation:

$$(22) \quad \ln(Q_{hprt}) = \alpha_p + \eta_p \ln(Y^R_{hrt}) + \vec{\beta}^{/}_p \vec{X}_{hrt} + \eta_p \varepsilon^{\bar{P}}_{prt} + \varepsilon_{hprt}$$

Unlike the earlier analysis, however, I now recognize that household income is not deterministically driven by regional real income and household educational attainment, but includes an error term, reflecting region x time specific variance:

$$(23) \quad \ln(Y^R_{hrt}) = \ln\left(Y^{R\sim E}_{rt}\right) + R_E E_{hrt} + \varepsilon^Y_{hrt} \quad \text{where} \quad \varepsilon^Y_{hrt} \sim N(0, \sigma^2_{Yrt})$$

The residual variation in incomes after accounting for education, $\varepsilon^Y$, produces correlation across the errors for the household product equations, so I now estimate all demand equations simultaneously using a household level random effects model:

$$(24) \quad \ln(Q_{hprt}) = \hat{a}_{prt} + \hat{b}_p R_E E_{hrt} + \vec{\hat{c}}^{/}_p X_{hrt} + \hat{b}_p e^Y_{hrt} + e_{hprt}$$

As before, all estimates are asymptotically consistent, with the product x region x time dummies ($\hat{a}$) capturing the product constant, the average regional education adjusted ln income, and the influence of regional relative prices:

$$(25) \quad \hat{b}_p = \eta_p \qquad \vec{\hat{c}}_p = \vec{\beta}_p \qquad \hat{\sigma}^2_{Yrt} = \sigma^2_{Yrt} \qquad \hat{a}_{prt} = \alpha_p + \eta_p \ln(Y^{R\sim E}_{rt}) + \eta_p \varepsilon^{\bar{P}}_{prt}$$

To calculate the product specific mean of the ln of regional and country real income, I add in the ln income implied by average regional educational attainment and then weight by the urban/rural household population shares:

$$(26) \quad E[\ln(\hat{Y}_{prt})] = \frac{\hat{a}_{prt}}{\hat{b}_p} + R_E \overline{\hat{E}}_{rt} \qquad E[\ln(\hat{Y}_{pct})] = \sum_{i=U,R} S_i E[\ln(\hat{Y}_{pit})]$$

To calculate the product specific ln of the mean, under the assumption of a ln normal distribution, I add in one-half of the residual income variance and, separately estimated, education variance:

$$(27) \quad \ln[E(\hat{Y}_{prt})] = \frac{\hat{a}_{prt}}{\hat{b}_p} + R_E \overline{\hat{E}}_{rt} + .5\hat{\sigma}^2_{Yrt} + .5R_E^2\hat{\sigma}^2_{Ert} \qquad \ln[E(\hat{Y}_{pct})] = \sum_{i=U,R} S_i e^{\ln[E(\hat{Y}_{pit})]}$$

These measures are then, as before, projected in a random effects panel regression on product dummies, growth rates, country level random effects and a variety of random effects designed to capture cross correlations in relative price levels and trends in product categories and groups

$$(28) \quad E[\ln(\hat{Y}_{pct})] \;\; or \;\; \ln[E(\hat{Y}_{pct})] = \hat{a}_p + \hat{g}t + u_c + e_p t + e_G t + e_{pc} + e_{Gc} + e_{pct}$$

producing estimates of the growth and cross sectional standard deviation of income, variously measured, under the assumption that these follow a global trend:

$$(29) \quad E[\ln(Y^R_{ct})] = E[\ln(Y^R_c)] + gt \quad or \quad \ln[E(Y^R_{ct})] = \ln[E(Y^R_c)] + gt$$

As before, it should be apparent that (28) can be run using urban/rural estimates alone (i.e. not the combined country value), their difference, or various versions of (26) or (27) excluding adjustments for educational attainment or within region dispersion, allowing for a more detailed analysis of the factors behind the levels and growth rates of income (see Section X). Furthermore, the first step estimates of the quasi-income elasticity and standard deviation of residual income ($\hat{b}_p$, $\hat{\sigma}_{Yrt}$) are inversed related to the return to education, $R_E$, so that the entire first step calculation can be run on educational attainment alone, with the coefficients subsequently modified by the preferred value of $R_E$. However, unlike the estimates of earlier sections, the estimate of the ln of the mean, (27), depends upon the square of the within region standard deviation of

educational attainment and residual income variation. Consequently, it is no longer possible to simply scale any of the estimated growth rates or cross-sectional variations of global income by the ratio of one's estimate of $R_E$ to my figure of .116. The estimated value of $R_E$ now affects the relative weight of trends in means versus trends in within region variation in calculating the overall trend of the ln of the mean, the standard measure reported in cross national data sets.[32]

[32] This section is clearly incomplete. First, the ln normal distribution is probably a poor approximation to income distribution (it most definitely is a poor approximation of the distribution of regional educational attainment). I am working on estimating more general and flexible income distribution functional forms to see whether the higher moments of these distributions affect the results. Second, in moving from the mean of the ln to the ln of the mean I have not yet made any adjustments for household size. Thus, at this time, I am moving my estimates from the mean of the ln household income to the ln of the mean household income, which is not exactly the ln of the mean per capita or equivalent adult income presented in cross national data sets.

## IX. Results III: The Growth, Level and Inequality of Real Consumption

Table XI below reports DHS based estimates of the growth and standard deviation of the mean of the ln and the ln of the mean of household consumption, estimated using the methods described in the previous section. As shown, moving to the standard national accounts based measure of living standards lowers the estimate of overall growth by about .6 percent, to 2.5 percent per annum.[33] This indicates a substantial downward trend in income inequality.[34] Incorporating dummies for the relative growth and level of sub-Saharan Africa, in the right hand side of the table, we see that this trend is concentrated in non sub-Saharan countries, where the ln of the mean growth rate is .9 percent lower, while sub-Saharan countries show a more modest trend, with the point estimate of growth only .3 percent lower.[35] After adjustment to the definitions at play in PWT, the estimated growth of non-African countries, at 2.3 percent, is within striking distance of the 1.7 to 1.9 percent indicated in the PWT data (Table X earlier). However, Africa's estimated growth rate, at 2.7 percent, remains as far away as ever from the PWT figure of .7 to .8 percent per annum.

Figures I and II below illustrate how the movement from the mean of the ln to the ln of the mean, i.e. the adjustment for income inequality, bring the DHS estimates closer to the PWT data. Figure I begins by graphing the estimates of country relative income levels in the year 2000

[33] The estimate of the mean and standard deviation of the mean of the ln is slightly different (i.e. somewhat higher) than that reported earlier in Section VI because, for this table, all of the demand equations are estimated, and linked, with the household random effect.

[34] As shown in a later draft, this stems from both a decline in urban-rural inequality and a decline in residual inequality (after adjustment for educational attainment) within rural and urban areas.

[35] These numbers appear different than the differences in the table because of the compounded errors introduced by rounding, i.e. the .023 is rounded up and the .031 is rounded down, while the African -.002 and .005 additions to .031 and .023 are both rounded up (in absolute terms).

| Table XI: DHS Estimates of the Mean of Ln versus Ln of the Mean<br>$y_{pct} = a_p + gt + e_pt + e_Gt + u_c + e_{pc} + e_{Gc} + e_{pct}$ | | | | |
|---|---|---|---|---|
| | Baseline | | African Levels & Growth | |
| | E[(lnY)] | ln[E(Y)] | E[(lnY)] | ln[E(Y)] |
| *g* | .031 (.005) | .025 (.006) | .031 (.005) | .023 (.006) |
| Africa's relative growth | | | -.002 (.003) | .005 (.003) |
| Africa's relative level | | | -.991 (.147) | -.899 (.142) |
| $\sigma_u$ | .711 (.073) | .668 (.069) | .510 (.056) | .494 (.054) |
| Note: Each specification includes the full set of error terms (p, G, pc, Gc, pct, etc) and product specific constants, but only g and $\sigma_u$ and the sub-Saharan related variables are reported. | | | | |

suggested by the two datasets.[36] Clearly, although the two sets of estimates often differ substantially and significantly,[37] the DHS numbers, whether estimated as the ln of the mean of the mean of the ln, are quite positively correlated with the PWT numbers. It is also apparent, however, that movement of the DHS concept to the ln of the mean moves the figures significantly closer to those of PWT. This is illustrated, further, in Figure II, where I graph the absolute change in the DHS estimate in moving from the mean of the ln to the ln of the mean against the initial deviation from PWT. Clearly, adjustment for the variation of income moves the two estimates systematically closer together. This shows, as argued earlier, that the concept implicit in the numbers reported by PWT and other crossnational data sets is a mixture of the average level of real living standards and its within country variation. I should note that the difference on the

[36] For the DHS, I run the specification of the two upper left-hand panels of Table XI, but with a fixed effects calculation of country dummies $u_c$. Similarly, the PWT data use the fixed effects version of the equation in Table III for ln consumption per equivalent adult. The data plotted in the figure are the demeaned fixed effects.

[37] A rough guide to the significance of differences can be taken from the fact that the average variance of the DHS and PWT country dummy coefficients is .043 or .041 (left & right panel) and .005 respectively, so that, roughly speaking, a standard deviation of .21 or .22 is associated with the difference between any two DHS and PWT estimates for a given country. The dashed lines drawn in the figure encompass the 95 percent region for this error bound.

# Figure I:  Real Consumption Measures

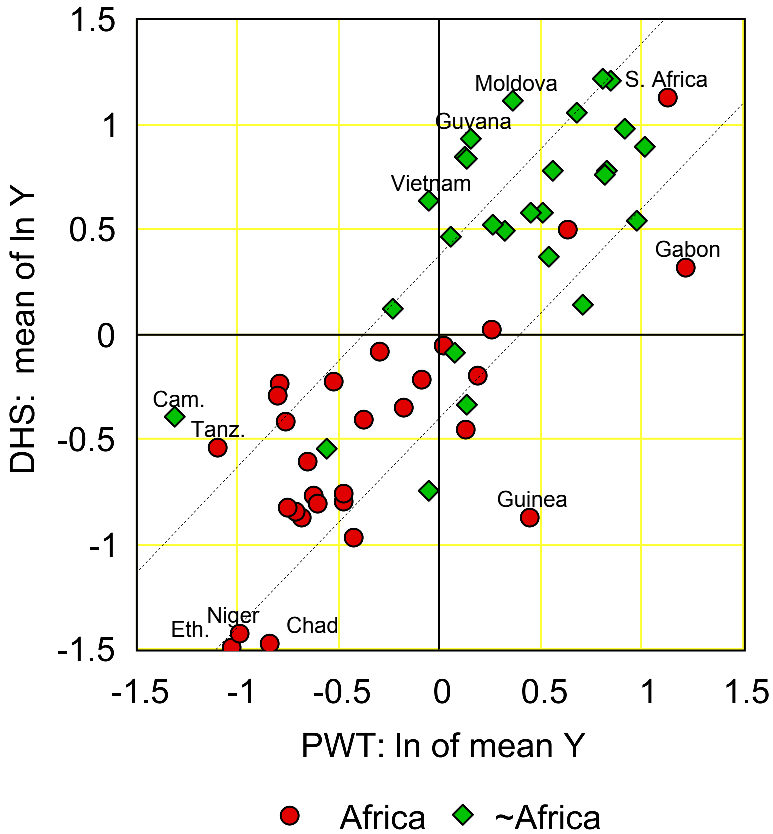


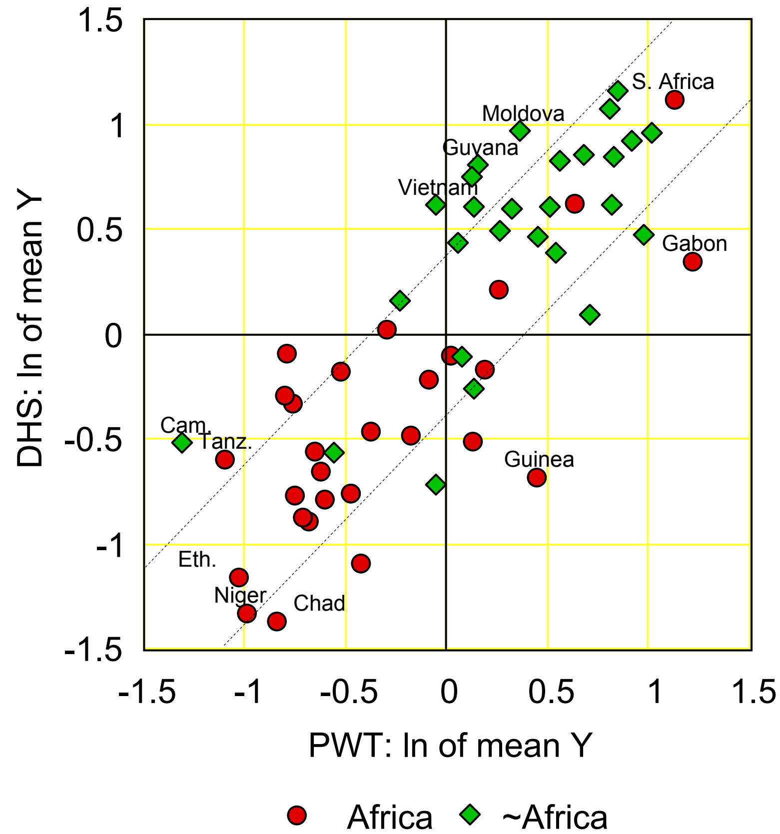

# Figure II:  Adjustment Toward PWT

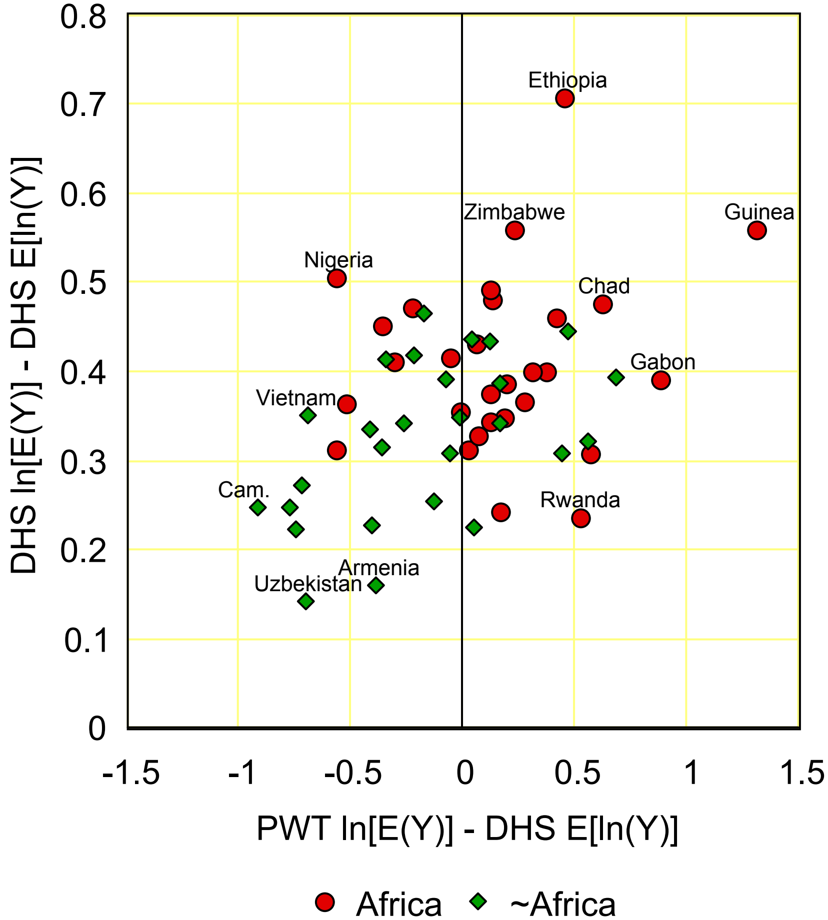

vertical axis of Figure II is an implicit measure of the within country variance of living standards and, as is readily apparent, indicates that levels of inequality are substantially higher in Africa. For this reason, the African level dummy, in the right hand panels of Table XI, is smaller (in absolute terms) with the ln of the mean concept, although it is still greater, albeit not significantly so, than that indicated by PWT (Table IX earlier).[38]

Figure III highlights the principal source of the discrepancy, in growth rates, between the DHS and PWT. The various panels project the country demeaned dummies (level estimates) for each country x survey time period against the country demeaned year.[39] This is the variation that identifies the growth rates estimated with each data source.[40] As shown in the upper two panels, the PWT and DHS ln of the mean data for the non sub-Saharan African countries are, in their overall pattern, quite consistent, indicating growth of about 2 percent per annum. Turning to the lower panel, however, one sees that sub-Saharan growth in PWT is negligible, while in the DHS data it is strong, clearly significant and on par with, or greater than, that present in non-African countries. It is extremely unlikely that any further adjustments of methodology can remove this. There is simply too much of an upward trend in the measured consumption of the DHS sub-

[38] The patterns shown in these figures are preliminary as I still have to modify the methodology to allow for non-normally distributed residual inequality and the adjustment from household to per capita figures, as noted in the previous section.

[39] For PWT, the country x survey time period dummy is simply the reported data. For the DHS, I run the specification $y_{pct} = a_p + u_{ct} + e_p*t + e_{pct}$ (the time trend and country x product random effect are no longer identified, and I still have to incorporate product group random effects into these estimates). In the figure, I remove the 14 countries with only one (time period) observation, as the residuals are automatically zero for both the DHS & PWT.

[40] The lines drawn in the panels are those implied by the African and non-African growth estimates of Table XI (DHS) and Table X (PWT). I should note that the growth regression, for the DHS, includes the first stage covariance and a random effects covariance matrix allowing for correlation across products in levels and trends. Consequently, the regression line is not simply the demeaned y variable projected against the demeaned x variable, but, as the reader can see, it is reasonably close to what that would be.

# Figure III:  Deviations of Levels from Country Means

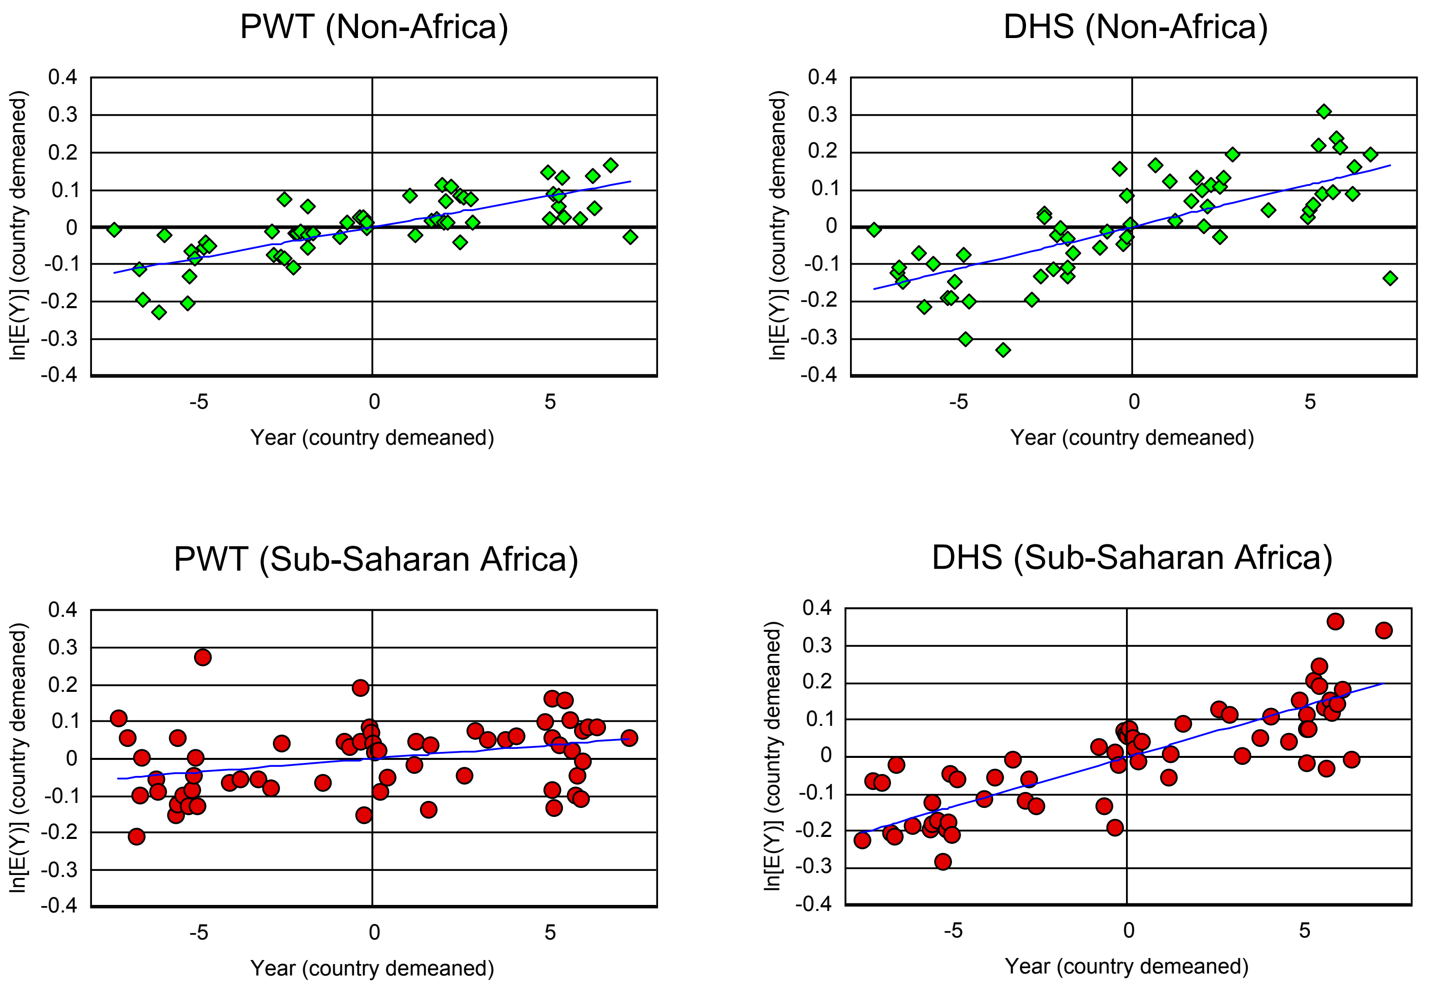

Saharan countries to be consistent with the utter stagnation implied by the PWT, and other cross national, data for the region.  African consumption, whether in the ln of the mean or the mean of the ln, is simply growing faster than cross national data sources, drawing on a mixture of country national accounts reports and ad hoc extrapolations and interpolations, indicate.

## X. Results IV: The Determinants of Living Standards and Inequality

Under construction.

# XI. Conclusion

Under construction. See the introduction for a summary of results.

## XII. Bibliography

Butler, J.S. and Robert Moffitt. "A Computationally Efficient Quadrature Procedure for the One Factor Multinomial Probit Model." Econometrica 50 (1982): 761-764.

Chamberlain, Gary. "Analysis of Covariance with Qualitative Data." Review of Economic Studies 47 (1980): 225-238.

Hall, Robert E. and Charles I. Jones. "Why Do Some Countries Produce So Much More Output per Worker than Others". Quarterly Journal of Economics 114 (1999): 83-116.

Klenow, Peter J. and Andres Rodriguez-Clare. "The Neoclassical Revival in Growth Economics: Has It Gone Too Far?" NBER Macroeconomics Annual 1997. B. Bernanke and J. Rotemberg, eds. Cambridge, MA: MIT Press, 73-102.

Psacharopoulos, George. "Returns to Investment in Education: A Global Update". World Development 22 (1994): 1325-1343.

# XIII. Appendix I: Demographic & Health Survey Data

Details on variables and cleaning up of DHS data in a later draft.

| DHS and Associated Surveys Used in the Paper | | | |
|---|---|---|---|
| Benin | 1996*, 2001, 2006 | Bolivia | 1994*, 1998*, 2003 |
| Burkina Faso | 1992, 1998, 2003 | Brazil | 1991, 1996 |
| Cameroon | 1991, 1998, 2004 | Colombia | 1990, 1995*, 2000, 2005 |
| Cen. Af. Rep. | 1994* | Dom. Rep. | 1991, 1996*, 1999, 2002 |
| Chad | 1996*, 2004 | Guatemala | 1995*, 1998* |
| Comoros | 1996* | Guyana | 2005 |
| Congo | 2005 | Haiti | 1994, 2000, 2005 |
| Cote D'Ivoire | 1994, 1998, 2005 | Honduras | 2005 |
| Ethiopia | 2000, 2005 | Nicaragua | 1997*, 2001 |
| Gabon | 2000 | Paraguay | 1990 |
| Ghana | 1993, 1998*, 2003 | Peru | 1992, 1996*, 2000, 2004 |
| Guinea | 1999, 2005 | | |
| Kenya | 1993, 1998, 2003 | Bangladesh | 1993, 1996, 1999, 2004 |
| Lesotho | 2004 | Cambodia | 2000, 2005 |
| Madagascar | 1992, 1997*, 2003 | India | 1992, 1998, 2005 |
| Malawi | 1992, 2000, 2004 | Indonesia | 1991, 1994, 1997, 2002 |
| Mali | 1995*, 2001, 2006 | Nepal | 1996*, 2001, 2006 |
| Mozambique | 1997*, 2003 | Pakistan | 1990 |
| Namibia | 1992, 2000 | Philippines | 1993, 1998*, 2003 |
| Niger | 1992, 1998, 2006 | Vietnam | 1997, 2002 |
| Nigeria | 1990, 1999*, 2003 | | |
| Rwanda | 1992, 2000, 2005 | Armenia | 2000, 2005 |
| Senegal | 1992, 2005 | Egypt | 1992, 1995*, 2000, 2003, 2005 |
| South Africa | 1998* | Kazakhstan | 1995, 1999 |
| Tanzania | 1992, 1996, 1999, 2003, 2004 | Kyrgyz Rep. | 1997 |
| Togo | 1998* | Moldova | 2005 |
| Uganda | 1995*, 2000, 2006 | Morocco | 1992, 2003 |
| Zambia | 1992, 1996*, 2001 | Turkey | 1993, 1998*, 2003 |
| Zimbabwe | 1994*, 1999, 2006 | Uzbekistan | 1996 |

Notes: Years denote date when survey began; data collection often continues into the following year. (*) Surveys with wage income data.